arXiv:1611.01688v1 [cs.LG] 5 Nov 2016

Oracle-Efficient Learning and Auction Design

Miroslav Dudík* Nika Haghtalab† Haipeng Luo‡ Robert E. Schapire§
Vasilis Syrgkanis¶ Jennifer Wortman Vaughan‖

November 8, 2016

Abstract

We consider the design of online learning algorithms in a general learning setting, with the goal of obtaining computationally efficient no-regret algorithms, when given access to an oracle for the offline optimization version of the problem. We present an algorithm that we call Generalized Follow-the-Perturbed-Leader (Generalized FTPL) and provide conditions under which it achieves vanishing regret and is efficiently implementable with an oracle. Our second main contribution is the introduction of a new online adversarial auction-design framework for revenue maximization and an application of our oracle-efficient learning results to the adaptive optimization of auctions.

Our learning algorithm is a generalization of the classic FTPL algorithm of Kalai and Vempala [27], playing at every iteration the historically best-performing action after adding some perturbation to the performance of each of its actions. The crux of our design is adding perturbations in a manner that leads to oracle-efficiency. We reduce this to designing a translation matrix, which translates a low dimensional vector with independent noise components into a high dimensional vector of perturbations on the learner's action space. Our approach generalizes prior work on oracle-efficient online learning [11, 22, 27, 35], ranging from online combinatorial optimization, learning in simultaneous auctions, and contextual learning.

Our auction-design framework considers an auctioneer learning an optimal auction rule in an online manner, every day observing an adversarially chosen vector of valuations. The auctioneer's goal is to achieve revenue that competes with the revenue of the optimal auction in hindsight among those in some target class. We give oracle-efficient learning results for: (1) VCG auctions with bidder-specific reserves in single-parameter settings with matroid constraints, (2) envy-free item pricing in multi-item auctions with unlimited supply, and (3) s-level auctions of Morgenstern and Roughgarden [28] for single-item settings. The last result implies good regret against the optimal overall auction when valuations are coming from a fast mixing Markov chain, that is independent across bidders. We also extend our results to the case when the learner observes side information on the bidders before running the auction (contextual learning).

We present additional extensions to contextual learning and learning with approximate oracles, implemented by FPTAS or Maximal-in-Range algorithms. We provide further applications in online welfare maximization in multi-unit auctions and in no-regret learning in simultaneous item auctions, answering an open question from prior work.

*Microsoft Research, New York, mdudik@microsoft.com
†Computer Science Department, Carnegie Mellon University, nhaghtal@cs.cmu.edu
‡Microsoft Research, New York, haipeng@microsoft.com
§Microsoft Research, New York, schapire@microsoft.com
¶Microsoft Research, New England, vasy@microsoft.com
‖Microsoft Research, New York, jenn@microsoft.com

1 Introduction

Online learning is increasingly playing a major role in the adaptive optimization of computer systems, from the design of online marketplaces [2, 5, 7, 33] to the optimization of routing schemes in communication networks [1]. The environments in these applications are constantly evolving in a non-stationary manner, requiring continued adaptation of these systems. Online learning robustly addresses this challenge by assuming that the environment is adversarial. The design of online learning algorithms for these settings has a long and distinguished history, starting from the seminal work of Freund and Schapire [15] and dating back to the very early work of Hannan [18] and Blackwell [4]. There are known learning rules that are information-theoretically optimal, with tight bounds on the improvement of solution quality as a function of the number of iterations of learning. However, the majority of online learning rules are computationally inefficient when the action space of the learner is exponential in the natural representation of the learning problem.

There is a line of research that addresses the design of computationally efficient online learning algorithms, primarily by reducing the online learning problem to an offline optimization problem, see e.g. [1, 22, 25, 27]. However, such reductions have found success only in limited albeit important settings, including online linear optimization, where the learner's action space is the power-set of some ground elements and the objective of the learner is a linear function of these elements [27]. However, a vast array of problems faced by online learning systems, particularly in online market design, is highly non-linear and more complex than online linear optimization.

Motivated by such applications, we address the design of *oracle-efficient online learning algorithms* in a general online learning setting with arbitrary learner objectives. Such algorithms are computationally efficient assuming access to an oracle for the corresponding offline optimization problem. From the recent work of Hazan and Koren [23], we know that, unlike the stochastic i.i.d. settings, oracle-efficient learning in adversarial environments is not achievable without further conditions on the learning problem. Hence, a primary contribution of our work is to give sufficient conditions on which problems are efficiently learnable online, and in particular, learnable via a generic algorithm that we introduce under the name *Generalized Follow-the-Perturbed-Leader (Generalized FTPL)*. Our algorithm is a significant generalization of prior approaches that worked only for specific settings, including the work of Kalai and Vempala [27] on online linear optimization, Hazan and Kale [22] on online submodular minimization, Daskalakis and Syrgkanis [11] on online learning in simultaneous second-price auctions, and Syrgkanis et al. [35] on adversarial contextual learning. Our sufficient conditions draw a strong connection between the notion of a universal identification set of Goldman et al. [16] and oracle-efficient learnability. Hence, our approach unifies and extends prior approaches to oracle-efficient learning.

The second main contribution of our work is to introduce a new framework for the problem of adaptive market design for revenue maximization and to demonstrate the power of our oracle-efficient algorithm by numerous applications in this framework. In our new *online optimal auction design* framework, at each iteration, a learner adaptively designs an auction rule for the allocation of a set of resources to a fresh set of players from a population.[1] The goal of the learner is to achieve average revenue at least as large as the revenue of the optimal auction from some target class. Unlike the standard approach to optimal auction design, initiated by the seminal work of Myerson [29], our approach is devoid of any assumptions about a prior distribution on the valuations of the players for the resources at sale. Instead, similar in nature to an agnostic approach in learning theory, we incorporate prior knowledge in the form of a target class of auction schemes that we want to compete with. A special case of our framework is the recent work of Roughgarden and Wang [33], which considers online learning over the class of single-item second-price auctions with bidder-specific reserves and gives an algorithm which achieves good regret with respect to a constant factor

[1] Equivalently the set of players at each iteration can be the same as long as players are not forward looking and only optimize their utility from the current iteration.

of the optimal revenue in hindsight. We go well beyond this specific setting and show that our *Generalized FTPL* algorithm can be used to optimize over several standard classes of auction rules which are known to approximate very well the optimal auctions in static stochastic environments. Examples include:

- the class of Vickrey-Clarkes-Groves (VCG) mechanisms with player-specific reserve prices in single-dimensional matroid auction settings, which are known to be a 2-approximation to the optimal mechanism in i.i.d. settings under some regularity conditions [21];
- the class of single-item s-level auctions introduced by Morgenstern and Roughgarden [28], who showed that these auctions approximate, to an arbitrary accuracy, the Myerson auction [29], which is known to be optimal for the Bayesian independent private value setting; and
- the class of envy-free item pricing mechanisms in combinatorial markets with unlimited supply, which has been widely studied in the static Bayesian setting [2, 17].

We present several extensions of our main results, including: (1) learning with side information (often referred to as *contextual* learning); (2) constant-factor approximate oracles, provided, e.g., by Maximal-in-Range (MIR) algorithms [30]; and (3) regret bounds with respect to stronger benchmarks for the case in which the environment is not completely adversarial but rather follows a fast mixing Markov process.

The contextual-learning extension is of particular interest to online auction design because it allows the learner to use any side information known about the bidders before they place their bids, to improve the auction. Although, typically, no two sets of bidders in the population are identical, the learner can utilize the side information to design a common treatment for bidders that are similar under some fixed representation, that is to *generalize across a population*. This is exactly achieved through the framework of contextual learning. While our optimality results, like much of the prior work, rely on the assumption that the bidders are either myopic or are different at each iteration, this assumption is in practice substantially less severe for contextual learning. The contextual learning algorithm can broaden the pool of bidders by including bidders from multiple markets, which reduces the probability that the exact same bidder will be used in the learning process multiple times and be overly influential in the choices of the algorithm. This in turn reduces ratchet effects where bidders identify that the auctioneer is learning and seek to strategically distort the information they provide so as to gain in the future.

The extension to the Markovian setting implies that under stochastic conditions on the process that generates the valuations of bidders, our online algorithm for s-level auctions has an average revenue that is at least as large as the revenue of the optimal Myerson auction for the stationary distribution of the Markov process, extending prior work that only gave such guarantees in i.i.d. settings [10, 12, 28, 32].

Finally, we provide further applications of our work in the context of online combinatorial optimization with MIR approximate oracles, and in the context of no-regret learning for bid optimization in simultaneous second-price auctions. In the former application, we give a polynomial time online learning algorithm for online welfare maximization in multi-unit auctions, that achieves 2-approximate regret by invoking the MIR approximation algorithm of Dobzinski and Nisan [13] as an offline oracle. With the latter application, we solve an open problem raised in the recent work of Daskalakis and Syrgkanis [11] who offered efficient learning algorithms only for the weaker benchmark of no-envy learning, rather than no-regret learning, in simultaneous second price auctions.

2 Overview of Main Results and Techniques

2.1 Oracle-Efficient Learning and the Generalized FTPL

We consider the general setting of online learning, where at each iteration $t \in \{1, \ldots, T\}$, the learner picks an action x_t from some finite set $\mathcal{X}$ and the adversary picks an action y_t from some set $\mathcal{Y}$. The learner receives a reward of $f(x_t, y_t)$. The goal of the learner is to choose an action at every iteration in such a way that his average reward is close to the reward of the best fixed action in hindsight, minimizing his average

regret:

$$\textsc{Average-Regret} = \max_{x \in \mathcal{X}} \frac{1}{T} \sum_{t=1}^{T} \left(f(x, y_t) - f(x_t, y_t) \right). \tag{1}$$

In particular, we seek algortihms for which this regret vanishes to zero as the time horizon T goes to infinity, at a rate that is polynomial in the representation of the learning problem and inversely polynomial with T.[2] Typically, the regret rate that we strive for depends only logarithmically on the size of the set of available learner actions $\mathcal{X}$.

Ignoring computational efficiency constraints, the exponential weights algorithm of Freund and Schapire [15] achieves a regret rate of $O(\sqrt{\log |\mathcal{X}|/T})$, which, absent any structure on the reward function, is information-theoretically optimal. However, this algorithm is computationally inefficient, as at every iteration, it updates a weight for each action in the learner's action space, thereby requiring time linear in $|\mathcal{X}|$. The action spaces we work with are frequently combinatorial, so the linear dependence in $|\mathcal{X}|$ is prohibitive. Can we achieve similar regret rates with computation that is polynomial in $\log |\mathcal{X}|$ and T?

Clearly this cannot be achieved without some restrictions on the reward structure. We adopt the common assumption that the learner has access to an oracle for the offline learning problem. Given a distribution over adversary actions represented as a uniform distribution over a set $S = \{y_1, \ldots, y_t\}$, the oracle finds the best learner action for the distribution, i.e.,

$$\textsc{Opt}(S) = \max_{x \in \mathcal{X}} \sum_{\tau=1}^{t} f(x, y_\tau) \tag{2}$$

The oracle is essentially solving the hindsight problem: given a history of adversary actions, find the best action for the learner in hindsight.

Why oracle-efficient learning? The assumption of an offline oracle is the minimal assumption needed to argue that the online problem is efficiently learnable. The offline problem (or, possibly, an FPTAS for the offline problem) is a weaker problem than the online learning problem, because standard online-to-batch reductions [6, 11] can turn a polynomial-time online learning algorithm into a polynomial-time additive-approximation scheme for the offline problem.[3] Hence, whatever assumptions are needed for an efficient online learning algorithm to exist, the same assumptions imply the existence of a computationally efficient oracle for the corresponding offline problem. Thus, under those assumptions, our algorithm would also be computationally efficient, implying that restricting the design of computationally efficient learning algorithms to oracle-efficient algorithms is without any loss of generality. An oracle-efficient algorithm can also be viewed as a reduction from the offline to the online problem, providing conditions under which the online problem is not only as hard, but also as easy as the offline problem, thereby offering a computational equivalence between online and offline optimization.

In addition to fundamental theoretical importance, oracle-efficient algorithms are advantageous from a practical viewpoint. For example, if one has already developed and implemented a Bayesian optimization procedure which optimizes against a static stochastic environment, then our algorithm offers a black-box transformation of that procedure into an adaptive optimization algorithm with provable learning guarantees in non-stationary, non-stochastic environments. Even if the existing Bayesian optimization system does not run in worst-case polynomial time, but is rather a well-performing fast heuristic, an oracle-efficient algorithm will leverage any expert domain knowledge that went into designing the heuristic to perform well for the problem at hand.

[2] For simplicity of exposition, we talk abstractly here, deferring a formal definition of what it means to be polynomial in the representation of the problem until later, because different learning problems have different parameters of interest.

[3] As we show in the appendix, our results extend under the weaker assumption of a polynomial-time additive-approximation scheme for the offline problem.

Generalized FTPL with Low-Dimensional Noise. With access to an oracle, a natural first attempt to design an online learner would be to use the Follow-the-Leader (FTL) algorithm. This algorithm simply invokes the oracle at every iteration and plays the best action in hindsight. In a stochastic environment, i.e., when the adversary's actions are drawn i.i.d. from some distribution at each iteration, such an approach works and offers regret rates on the order $O(\log |\mathcal{X}|/\sqrt{T})$ by a simple application of the Rademacher complexity of finite hypotheses classes [34]. However, the lack of randomness in this algorithm leads to very high regret in adversarial environments. A seminal paper of Kalai and Vempala [27] showed that a slight modification of this algorithm leads to good regret. At a high level, this modification perturbs the performance of each action $x \in \mathcal{X}$ by some random amount ϵ_x drawn from a high-variance distribution, before choosing the best action in hindsight. However, even though this algorithm has good regret, it is computationally inefficient as it requires putting a perturbation on the performance of each one of the actions available to the learner. Moreover, this algorithm does not leverage the existence of the offline oracle, as the perturbed optimization problem is not a problem that the oracle is able to solve.

Daskalakis and Syrgkanis [11] introduced the approach of adding randomness to the FTL algorithm by augmenting the observed history with "fake" historical samples of adversary actions. The oracle is then invoked on the resulting perturbed history of adversary actions. However, they did not give any general way of generating such perturbations of the history or any general conditions under which this approach leads to good regret or computational efficiency when the set of adversary's actions is large. Our work exactly fills in this missing piece.

Perturbation Translation Matrices. Our results bridge the gap between these two approaches. We construct perturbations on each individual action of the learner by *coupling* the randomness across different actions. We generate a small number N of independent perturbations $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_N)$ drawn from a highly dispersed distribution. In most applications, N is logarithmic in $|\mathcal{X}|$. We then define the perturbation on the performance of each learner action x to be a weighted combination of these independent perturbations. In this way, we construct a *perturbation translation matrix* $\boldsymbol{\Gamma}$ (for simplicity assume entries in $[0, 1]$), which translates the short N-dimensional perturbation vector into a long $|\mathcal{X}|$-dimensional one, i.e., the perturbation on action x is $\epsilon_x = \boldsymbol{\alpha} \cdot \boldsymbol{\Gamma}_x$. By sampling the low dimensional noise vector $\boldsymbol{\alpha}$, we avoid the exponential computation of drawing a separate perturbation term for each action of the learner, while at the same time exploiting the idea that high variance noise added on each action of the learner leads to good regret. We term this algorithm *Generalized FTPL.*

Generalized FTPL has low regret if the translation matrix $\boldsymbol{\Gamma}$ satisfies an *admissibility* criterion: the main property required from $\boldsymbol{\Gamma}$ is that its rows are (sufficiently) distinct. Thus each action's perturbation uses a different weighted combination of the low dimensional noise. We show that any such matrix leads to vanishing regret, with regret rates that are more concretely captured in the following theorem (see Theorem 3.5 for a formal version):

Informal Theorem 1. *A translation matrix is (κ, δ)-admissible if any two rows of the matrix are distinct, the number of different values within a column is at most κ and the minimum non-zero difference between any two values within a column is at least δ. Generalized FTPL with a (κ, δ)-admissible matrix achieves regret of $O(N\sqrt{T\kappa/\delta})$.*

Oracle-based Implementability. To complete the construction of an oracle-efficient algorithm, we need the property that these induced action-level perturbations can be simulated by a perturbation of the history, i.e., that there exists a distribution of adversary actions that induces the same randomness as the one induced by $\epsilon_x = \boldsymbol{\alpha} \cdot \boldsymbol{\Gamma}_x$. This is captured by the *implementability* criterion, which states that each column of the translation matrix essentially corresponds to a scaled version of the expected reward of the learner on some distribution of adversary actions (see Theorem 3.11 for a formal statement).

Informal Theorem 2. *A translation matrix is implementable if each column corresponds to a scaled version of the expected reward of the learner against some small-supported distribution of actions of the adversary. Generalized FTPL with an implementable and (κ, δ)-admissible translation matrix can be implemented in an oracle-efficient manner, with running time polynomial in N, T, κ/δ, and in the size of the distribution implementing the translation matrix, when represented by a sequence of atoms.*

Design Approach. Thus we get that a sufficient condition for oracle-efficient learnability is the existence of an admissible and implementable translation matrix, with small N, κ and $1/\delta$. For some learning problems, it is easier to construct a trivially implementable translation matrix and then argue about its admissibility; for others, it is easier to start with a trivially admissible matrix and argue about its implementability. Below we describe examples of each, exhibiting the versatility of our conditions.

The first case, when implementability is immediate, is when each column corresponds to the reward of the learner on some particular adversary action. The admissibility translates to a natural condition: the reward of the adversary under any two actions $x, x' \in \mathcal{X}$ must be different for at least one of the N adversary actions that are associated with the columns of $\mathbf{\Gamma}$. Therefore, to construct an admissible and implementable matrix $\mathbf{\Gamma}$, it suffices to identify a small number N of adversary actions that have this differentiating property.

Corollary 1. *If there exists a small set of adversary actions such that any two learner's actions yield a different reward on at least one of these adversary actions and such that the number of different reward levels for each of these actions in the set is small, then the learning problem is oracle-efficient and learnable with the Generalized FTPL algorithm.*

A special case in which this holds is captured by the notion of a *universal identification sequence* defined by Goldman et al. [16]. This notion applies when the learner is choosing binary functions from some input space and the adversary is choosing realizations of the input variables, with the learner collecting a reward only if the output of his function on the adversary's input is 1. Another special case of this differentiating property arises in contextual learning. In contextual learning, the learner chooses a policy that maps a domain of contextual information into one of K actions, with the adversary picking both the contextual information and the reward of each action. The learner needs to compete with the optimal policy in a given policy class. A *separator* [11] is a small set of contexts such that any two policies in the policy class choose different actions on at least one context from the set. The existence of a separator implies an admissible and implementable translation matrix.

Instead of beginning the design with implementability, we can begin with admissibility. One interesting special case of an admissible matrix is when the rows of matrix $\mathbf{\Gamma}$ contain a binary representation of some parameterization of the actions of the learner. The number of columns of the matrix is then logarithmic in the number of actions of the learner. However, such a matrix may or may not be implementable.

Corollary 2. *If any linear functions on some binary vector representation of the learner's actions, can be implemented with a distribution of adversary actions, then the learning problem is oracle-efficient and learnable with the Generalized FTPL algorithm.*

This condition is perhaps easiest to understand when the learner's actions are a power-set over a set of K ground elements, i.e., $\mathcal{X} = \{0, 1\}^K$. This is the case for online combinatorial optimization in which the adversary picks a combinatorial function v_t at every iteration and the learner receives reward $v_t(x_t)$. The perturbations of our algorithm essentially correspond to adding an independent perturbation on each of the ground elements, or equivalently, adding a fake additive function to the history, whose value for each ground element is drawn independently from some distribution. If the adversary is picking functions from a class that includes additive ones, then the latter perturbations are by definition implementable. This approach was leveraged by Hazan and Kale [22] for online submodular minimization. Their approach is a special case of ours, but our work extends beyond combinatorial optimization. For instance, in subsequent sections,

we show that a translation matrix consisting of binary encodings of reserve prices in auctions yields an oracle-efficient online learning algorithm for revenue maximization.

2.2 Main Application: Online Optimal Auction Design

As a main application of our oracle-efficient learning algorithm, we consider an online learning version of the classical market design problem of optimally selling a single item (or multiple items) to a set of interested buyers so as to maximize revenue. Traditional optimal auction theory assumes that the valuations of the players are drawn from a population distribution which is known, thereby leading to a Bayesian optimization problem. The knowledge of the distribution by the seller is a rather harsh assumption and recent work in algorithmic mechanism design [10, 12, 28, 32] relaxes this assumption by solely assuming access to a set of samples from the distribution. In this work we take a completely distribution-free and adversarial approach, introducing the setting of *online optimal auction design*.

In many applications of optimal auction theory, in particular in electronic marketplaces, the seller does not sell the item once, but rather repeatedly sells the item to a population of buyers, with a few arriving at each auction. Hence, the seller can optimize his auction design in an online manner, using historical data consisting of observed valuations. The goal of the seller is to leverage the historical data and pick an auction at each iteration so as to compete with the optimal auction from some target class. For instance, in a sponsored-search auction setting, the auctioneer can use historical data from past auctions in a sub-market to decide what reserve price to set for a new impression that arrives in that sub-market. In a setting like eBay, the system can use historical bid data from auctions of similar goods to decide what reserve price to recommend to a new seller.

Motivated by such online auction optimization settings we consider the following optimal auction model. An auctioneer is repeatedly selling a set of resources for T time-steps. On each day t a set of n bidders arrive with a vector of valuations/bids $\mathbf{v}_t \in \mathcal{V}^n$. Prior to observing the bids, the auctioneer commits to an auction a_t from a class of truthful auctions $\mathcal{A}$ that he is interested in. Unlike classical optimal auction design, we make no assumptions about the process which generates the valuations of players on each day. We consider the case where these valuations are chosen by an adversary.

The goal of the learner is to achieve an average revenue over time, that, in hindsight, is very close to the revenue that would have been achieved by the best fixed mechanism in class $\mathcal{A}$, if that mechanism was at all time-steps, i.e., the learner aims to achieve small average regret:

$$\textsc{Average-Regret} = \max_{a \in \mathcal{A}} \frac{1}{T} \sum_{t=1}^{T} \mathrm{Rev}(a, \mathbf{v}_t) - \frac{1}{T} \sum_{t=1}^{T} \mathrm{Rev}(a_t, \mathbf{v}_t). \tag{3}$$

We apply our oracle-efficient algorithm and design online learning algorithms for the aforementioned setting assuming access to an offline oracle. In the context of revenue maximization, the learner's action space is the set of target auctions $\mathcal{A}$, the adversary's action space is the set of valuation vectors $\mathcal{V}^n$. Finally, the offline oracle is a Bayesian revenue maximization oracle which computes the optimal auction within a class of auctions $\mathcal{A}$, for any correlated distribution of valuations given as a uniform distribution over a set of valuation vectors: $\{\mathbf{v}_1, \ldots, \mathbf{v}_k\}$.

We construct oracle-efficient online learning algorithms for three optimal auction classes, mentioned in Section 1, which have been well studied in the Bayesian optimal auction literature. For each of these auction settings, we show how to construct a translation matrix that is admissible and implementable.

Example: Second-Price with Bidder-Specific Reserves. To give a hint at our techniques we briefly describe our construction for one very special case of our online optimal auction design, the case of second-price auctions with bidder-specific reserves, in which the auctioneer ignores bidders below their reserve price and allocates to the highest remaining bidder, charging him the maximum of his reserve and the

second highest remaining bid. Thus the learner needs to pick at each iteration a reserve price for each bidder. For simplicity assume that each reserve price comes from a discretized grid of $m+1$ possible levels $R = \{0, 1/m, 2/m, \ldots, 1\}$ and valuations are upper bounded by 1. Thus the action space of the learner is of size $(m+1)^n$, which is prohibitively large. Our construction yields an oracle-efficient algorithm which runs in time $\text{poly}(n, m, T)$, ignoring oracle computation, and yields regret $O(n \log(m)/\sqrt{T})$.

We present here a simpler translation matrix which yields regret $O(n \cdot m^2/\sqrt{T})$. We use the approach implied by our Corollary 1 and find a set of special bid profiles such that any two auctions in the class achieve different revenue in at least one of these bid profiles. It is easy to find such a set: for each bidder i, consider all the bid profiles where he bids some level $h \in R$ and every other bidder bids 0. For any two different auctions in the class of second-price auctions with a reserve, the revenue must be different in at least one of these bids: if the auctions differ on the reserve price of bidder i, with one auction having reserve r and the other having reserve r', then a bid profile which submits a bid level $h \in (r, r']$ for bidder i and a zero bid on everyone else, will yield different revenue in the two auctions. Thus we have constructed a translation matrix where each column corresponds to the revenue of the auction under the associated bid from one of these $N = n \cdot (m+1)$ bids. Moreover, observe that for each of these special bid profiles the revenue of any auction takes at most $m+1$ different values and the minimum non-zero difference in the revenue of two auctions is $1/m$. Thus invoking the Informal Theorems 1 and 2 we get that there exists an oracle-efficient no-regret algorithm which runs in time $poly(n, m, T)$, ignoring oracle computation, and achieves average regret $O(nm^2/\sqrt{T})$.

Markovian adversaries and competing with overall optimal. Another class of auctions that we analyze is the recently introduced class of *level auctions* [28]. A level auction is a generalization of bidder-specific reserves, where each bidder now has multiple associated reserves which categorize his bid into a set of buckets that are then used for allocating the good. Morgenstern and Roughgarden [28] showed that these auctions, for large enough number of buckets, provide an arbitrarily accurate approximation to the overall optimal Myerson auction in the Bayesian single-item auction setting, where the value of each bidder is drawn from an independent distribution. Our learning results imply oracle-efficient regret bounds of $O(n^2 m^2/\sqrt{T})$. If the valuation of each player is drawn independently in each round from some fixed distribution, then standard online-to-batch reduction implies that the revenue of the online learning algorithm is close to the overall optimal single-shot auction, i.e., the Myerson auction. We generalize this reasoning and show that we get such strong optimality guarantee also when the valuations of players are drawn at each iteration from a fast-mixing Markov process, which is independent across players. For both of these setting, our results give an oracle-efficient algorithm with regret $O(n^{2/3}/T^{1/6})$ to the overall optimal (by optimizing over the value of m). This is the first result on competing with a Myerson optimal auction for non-iid distributions, unlike prior work [10, 12, 28, 32] which assumes i.i.d. samples.

3 Generalized FTPL and Oracle-Efficient Online Learning

In this section, we introduce a general Follow-the-Perturbed-Leader (FTPL) algorithm. We then describe the conditions under which this algorithm efficiently reduces online learning to offline optimization.

We consider the following online learning problem. On each round $t = 1, \ldots, T$, a learner chooses an action x_t from a finite set $\mathcal{X}$, and an adversary chooses an action y_t from set $\mathcal{Y}$, which is not necessarily finite. The learner then observes y_t and receives a payoff $f(x_t, y_t) \in [0, 1]$, where the function f is fixed and known to the learner. The goal of the learner is to obtain low expected regret with respect to the best action in hindsight, i.e., to minimize

$$\textsc{Regret} := \mathbb{E}\left[\max_{x \in \mathcal{X}} \sum_{t=1}^{T} f(x, y_t) - \sum_{t=1}^{T} f(x_t, y_t)\right],$$

where the expectation is over the randomness of the learner. To simplify exposition, we assume that the ad-

versary is oblivious, i.e., that the sequence $y_1, \ldots, y_T$ is chosen up front without knowledge of the learner's realized actions. Our results generalize to adaptive adversaries using standard techniques [11, 24].

Our algorithm broadly follows the FTPL scheme of Kalai and Vempala [27]. In the original FTPL scheme of Kalai and Vempala [27], at every iteration the algorithm adds a perturbation to the historical cumulative payoff of each possible action of the learner and then picks the action with the largest perturbed historical payoff. The perturbation on each action is independently drawn from a highly dispersed distribution, such as an exponential distribution with standard deviation of order $\sqrt{T}$. Applying this approach to the general learning setting above is computationally inefficient when the action space of the learner $\mathcal{X}$ is exponential in some succinct representation of the learning setting.

For online combinatorial optimization with either additive [27] or submodular [22] objectives, it has been shown that one does not need to add a perturbation on each set, but adding a perturbation on each item suffices to lead to no-regret. However, this approach is specific to these settings and it was not even known before our work whether the approach extends to online combinatorial optimization with arbitrary objectives. We show how to bypass the exponential perturbation problem for the general online learning setting.

Our goal is to create perturbations that can be succinctly passed on to an offline optimization oracle. In particular, we seek to generate perturbations that can be simulated by augmenting the history with a polynomial number of fake samples as proposed in Daskalakis and Syrgkanis [11]. To generate such perturbations, we begin by drawing a random vector $\boldsymbol{\alpha} \in \mathbb{R}^N$ of some small size N, with components α_j drawn independently from some dispersed distribution D. The payoff of each of the learner's action is perturbed by a linear combination of these independent variables, as prescribed by a non-negative translation matrix $\boldsymbol{\Gamma}$ of size $|\mathcal{X}| \times N$, i.e., the vector of perturbations on all the $|\mathcal{X}|$ actions of the learner is given by the vector $\boldsymbol{\Gamma}\boldsymbol{\alpha}$. Letting $\boldsymbol{\Gamma}_x$ denote the row of $\boldsymbol{\Gamma}$ corresponding to x, on each round t, the algorithm outputs the action x that maximizes $\sum_{\tau=1}^{t-1} f(x, y_\tau) + \boldsymbol{\alpha} \cdot \boldsymbol{\Gamma}_x$. See Algorithm 1 for a full specification. For non-oblivious adversaries, a fresh random vector $\boldsymbol{\alpha}$ is drawn in each round.

In this section we analyze the properties of the matrix $\boldsymbol{\Gamma}$ to guarantee that the Generalized FTPL algorithm both achieves a small regret and its perturbations can be efficiently transformed into fake samples. Together these properties give rise to efficient reductions of online learning to offline optimization.

Algorithm 1: Generalized FTPL

Input: non-negative matrix $\boldsymbol{\Gamma} \in \mathbb{R}_+^{|\mathcal{X}| \times N}$, distribution D
Draw $\alpha_j \sim D$ for $j = 1, \ldots, N$.
for $t = 1, \ldots, T$ **do**
  Choose action $x_t = \arg\max_{x \in \mathcal{X}} \left[\sum_{\tau=1}^{t-1} f(x, y_\tau) + \boldsymbol{\alpha} \cdot \boldsymbol{\Gamma}_x \right]$.
  Observe y_t and receive payoff $f(x_t, y_t)$.
end for

3.1 Regret Analysis

To analyze the Generalized FTPL algorithm, we first bound its regret by the sum of a *stability* term and *perturbation* term in the following lemma. While this approach is standard [27], we include a proof in Appendix A for completeness.

Lemma 3.1 (FTPL Lemma)**.** *For the Generalized FTPL algorithm,*

$$\textsc{Regret} \leq \mathbb{E}\left[\sum_{t=1}^{T} f(x_{t+1}, y_t) - f(x_t, y_t) \right] + \mathbb{E}\left[\boldsymbol{\alpha} \cdot (\boldsymbol{\Gamma}_{x_1} - \boldsymbol{\Gamma}_{x^*}) \right] \tag{4}$$

where $x^* = \arg\max_{x \in \mathcal{X}} \sum_{t=1}^{T} f(x, y_t)$.

The first term measures the stability: how much our decision changes from round to round. The second term measures the strength of perturbation: how much the perturbation amount differs between the best action and the initial action. To bound the stability term, we require that the matrix $\mathbf{\Gamma}$ be *admissible* and the distribution D be *dispersed* in the following sense:

Definition 3.2 ((κ, δ)-Admissible Translation Matrix)**.** *A translation matrix* $\mathbf{\Gamma}$ *is* admissible *if its rows are distinct. It is* (κ, δ)-admissible *if it is admissible and also:*

1. *the number of distinct elements within each column is at most* κ*,*
2. *distinct elements within each column differ by at least* δ*.*

Definition 3.3 ((ρ, L)-Dispersed Distribution)**.** *A distribution* D *on the real line is* (ρ, L)*-dispersed if for any interval of length* L*, the probability measure placed by* D *on this interval is at most* ρ*.*

In the next lemma, we bound the stability term in Equation (4) by showing that with high probability, for all rounds t, we have $x_{t+1} = x_t$. Since all rows of an admissible matrix $\mathbf{\Gamma}$ are distinct, it suffices to show that the probability that $\mathbf{\Gamma}_{x_{t+1}} \neq \mathbf{\Gamma}_{x_t}$ is small. We prove this for each coordinate $\Gamma_{x_{t+1}j}$ separately, by showing that it is only possible to have $\Gamma_{x_{t+1}j} \neq \Gamma_{x_t j}$ when the random variable α_j falls in a small interval, which happens with only small probability for a sufficiently dispersed distribution D. Note that any admissible matrix is (κ, δ)-admissible for some choices of κ and δ. The magnitudes of κ and δ and the number of columns of $\mathbf{\Gamma}$ determine how large the stability term can get.

Lemma 3.4. *Consider the Generalized FTPL algorithm with a* (κ, δ)*-admissible matrix* $\mathbf{\Gamma}$ *with* N *columns and a* $(\rho, 1/\delta)$*-dispersed distribution* D*. Then,*

$$\mathbb{E}\left[\sum_{t=1}^{T} f(x_{t+1}, y_t) - f(x_t, y_t)\right] \leq 2TN\kappa\rho.$$

Proof. Fix any $t \leq T$. The bulk of the proof will establish that, with high probability, $\mathbf{\Gamma}_{x_{t+1}} = \mathbf{\Gamma}_{x_t}$, which by admissibility implies that $x_{t+1} = x_t$ and therefore $f(x_{t+1}, y_t) - f(x_t, y_t) = 0$.

Fix any $j \leq N$. We first show that $\Gamma_{x_{t+1}j} = \Gamma_{x_t j}$ with high probability. Let V denote the set of values that appear in the j^{th} column of $\mathbf{\Gamma}$. For any value $v \in V$, let x^v be the action that maximizes the perturbed profit, ignoring the perturbation of column j itself, among those whose $\mathbf{\Gamma}$ entry in the j^{th} column equals v:

$$x^v \coloneqq \operatorname*{arg\,max}_{x \in \mathcal{X}:\, \Gamma_{xj} = v} \left[\sum_{\tau=1}^{t-1} f(x, y_\tau) + \sum_{j' \neq j} \alpha_{j'} \Gamma_{xj'}\right] = \operatorname*{arg\,max}_{x \in \mathcal{X}:\, \Gamma_{xj} = v} \left[\sum_{\tau=1}^{t-1} f(x, y_\tau) + \boldsymbol{\alpha} \cdot \mathbf{\Gamma}_x - \alpha_j v\right].$$

For any $v, v' \in V$, define

$$\Delta_{vv'} = \left(\sum_{\tau=1}^{t-1} f(x^v, y_\tau) + \boldsymbol{\alpha} \cdot \mathbf{\Gamma}_{x^v} - \alpha_j v\right) - \left(\sum_{\tau=1}^{t-1} f(x^{v'}, y_\tau) + \boldsymbol{\alpha} \cdot \mathbf{\Gamma}_{x^{v'}} - \alpha_j v'\right).$$

Note that x^v and $\Delta_{vv'}$ are independent of α_j, as we removed the payoff perturbation corresponding to α_j.

If $\Gamma_{x_t j} = v$, then by the optimality of x_t on the perturbed profit, we have $\alpha_j(v' - v) \leq \Delta_{vv'}$ for all $v' \in V$. Suppose $\Gamma_{x_{t+1}j} \neq v$. Then there is some $v' \in V$ which yields a better perturbed profit in the next round, i.e.,

$$\sum_{\tau=1}^{t-1} f(x^{v'}, y_\tau) + f(x^{v'}, y_t) + \boldsymbol{\alpha} \cdot \mathbf{\Gamma}_{x^{v'}} \geq \sum_{\tau=1}^{t-1} f(x^v, y_\tau) + f(x^v, y_t) + \boldsymbol{\alpha} \cdot \mathbf{\Gamma}_{x^v}.$$

Rearranging, we obtain for this same v' that

$$\Delta_{vv'} \le \alpha_j(v'-v) + f(x^{v'}, y_t) - f(x^v, y_t) \le \alpha_j(v'-v) + 1.$$

If $v' > v$, then

$$\alpha_j \ge \frac{\Delta_{vv'} - 1}{v' - v} \ge \min_{\hat{v} \in V,\, \hat{v} > v} \frac{\Delta_{v\hat{v}} - 1}{\hat{v} - v}$$

and so $\alpha_j(\overline{v} - v) + 1 \ge \Delta_{v\overline{v}}$ where $\overline{v}$ is the value of $\hat{v}$ minimizing the expression on the right. Thus, in this case we have $0 \le \Delta_{v\overline{v}} - \alpha_j(\overline{v} - v) \le 1$. Similarly, if $v' < v$, then

$$\alpha_j \le \frac{\Delta_{vv'} - 1}{v' - v} \le \max_{\hat{v} \in V,\, \hat{v} < v} \frac{\Delta_{v\hat{v}} - 1}{\hat{v} - v}$$

and so $\alpha_j(\underline{v} - v) + 1 \ge \Delta_{v\underline{v}}$ where $\underline{v}$ is the value maximizing the expression on the right. In this case we have $0 \le \Delta_{v\underline{v}} - \alpha_j(\underline{v} - v) \le 1$. Putting this all together, we have

$$\begin{aligned}
&\Pr\left[\Gamma_{x_{t+1}j} \ne \Gamma_{x_t j} \,\middle|\, \alpha_k, k \ne j\right] \\
&\quad\le \Pr\left[\exists v \in V:\ 0 \le \Delta_{v\overline{v}} - \alpha_j(\overline{v} - v) \le 1 \text{ or } 0 \le \Delta_{v\underline{v}} - \alpha_j(\underline{v} - v) \le 1 \,\middle|\, \alpha_k, k \ne j\right] \\
&\quad\le \sum_{v \in V} \left(\Pr\left[\alpha_j \in \left[\tfrac{\Delta_{v\overline{v}}-1}{\overline{v}-v}, \tfrac{\Delta_{v\overline{v}}}{\overline{v}-v}\right] \,\middle|\, \alpha_k, k \ne j\right] + \Pr\left[\alpha_j \in \left[\tfrac{-\Delta_{v\underline{v}}}{v-\underline{v}}, \tfrac{-\Delta_{v\underline{v}}+1}{v-\underline{v}}\right] \,\middle|\, \alpha_k, k \ne j\right] \right) \\
&\quad\le 2\kappa\rho.
\end{aligned}$$

The last line follows from the fact that $\overline{v} - v \ge \delta$ and $v - \underline{v} \ge \delta$, the fact that D is $(\rho, 1/\delta)$-dispersed, and a union bound.

Since this bound does not depend on the values of the α_j, we can remove the conditioning and bound $\Pr[\Gamma_{x_{t+1}j} \ne \Gamma_{x_t j}] \le 2\kappa\rho$. Taking a union bound over all $j \le N$, we then have that, by admissibility, $\Pr[x_{t+1} \ne x_t] = \Pr\left[\mathbf{\Gamma}_{x_{t+1}} \ne \mathbf{\Gamma}_{x_t}\right] \le 2N\kappa\rho$, which implies the result. □

To bound the regret, it remains to bound the perturbation term in Equation (4). This bound is specific to the distribution D. Many distribution families, including uniform, Gaussian, Laplacian, discrete uniform, exponential etc., with appropriately set variance, will lead to a sublinear regret. Here we present a concrete regret analysis for the case of a uniform distribution supported on $[0, 1/\eta]$, for $\eta \approx 1/\sqrt{T}$:

Theorem 3.5. *Let $\mathbf{\Gamma}$ be a (κ, δ)-admissible matrix with N columns and $\gamma = \max_{x,j} \Gamma_{xj}$ and let D be the uniform distribution on $[0, 1/\eta]$ for $\eta = \sqrt{\gamma\delta/(2T\kappa)}$. Then the regret of the Generalized FTPL algorithm can be bounded as*

$$\textsc{Regret} \le 2N\sqrt{2T\kappa\gamma/\delta}.$$

The proof of Theorem 3.5 follows immediately from Lemmas 3.1 and 3.4 with $\rho = \eta/\delta$.

3.2 Oracle-Efficient Online Learning

The Generalized FTPL algorithm requires the ability to optimize the perturbed objective over the learner's action set in order to find x_t at each round t. However, to achieve online-to-offline reduction, we need to create offline optimization problems of the same type as the underlying online problem. Namely, we need to reduce to optimization of the cumulative payoff over some data set consisting of adversary's actions. An algorithm solving such offline optimization problems will be called an *offline oracle*. If each optimization step of the Generalized FTPL can be efficiently transformed into a sequence of oracle calls, we will say that the Generalized FTPL algorithm is *oracle-efficient*.

In this section, we derive the criterion that allows us to solve the optimization of the perturbed objective by a single oracle call. We achieve this by transforming the perturbation into a (weighted) set of adversary's actions, which are then passed to the oracle alongside all of the actual adversary's actions.

As a result of this construction, we immediately get that the online learning problem can be solved efficiently whenever there is an efficient offline oracle and the Generalized FTPL is oracle-efficient.

We now define the offline oracle and oracle-efficient online learning more formally.

Definition 3.6 (Offline Oracle)**.** *An* offline oracle OPT *is any algorithm that receives as input a weighted set of adversary actions $S = \{(w_\ell, y_\ell)\}_{\ell \in \mathcal{L}}$, $w_k \in \mathbb{R}^+$, $y_k \in \mathcal{Y}$, and returns*

$$\textsc{Opt}(S) \in \arg\max_{x \in \mathcal{X}} \sum_{(w,y) \in S} w f(x, y).$$

Definition 3.7 (Mass of a Data Set)**.** *The* mass *of a weighted dataset S, denoted $\|S\|$, is the maximum of its cardinality and its sum of weights, $\|S\| \coloneqq \max\left\{|S|, \sum_{(w,y)\in S} w\right\}$.*

Definition 3.8 (Oracle-Efficient Online Learning)**.** *We say that an online learning algorithm is* oracle-efficient *with the per-round complexity g(T) if in each round it makes $\mathcal{O}(g(T))$ calls to an offline oracle, each on a dataset with the mass $\mathcal{O}(g(T))$, and its per-round running time, excluding oracle calls, is $\mathcal{O}(g(T))$.*

Note that unlike the oracle of [23], which accepts a distribution, we do not require the weights to sum to one. This is without loss of generality since as we show in Appendix C the two oracle classes are essentially equivalent, as we show that any oracle efficient algorithm that uses a fractionally weighted oracle can be implemented with an integral oracle with an arbitrarily small increase in regret. For now, we mention that the reason why we introduced the mass of a dataset in the definition of oracle-efficiency is that we anticipate reductions to integral oracles, which typically expect unweighted inputs, obtained by replicating each example according to its weight. Thus the mass of a weighted data set is the natural problem size for integral oracles.

We next provide the criterion on the pair $(\mathbf{\Gamma}, D)$, which allows us to transform the perturbed objective into a weighted dataset, thus achieving oracle-efficiency of the Generalized FTPL:

Definition 3.9. *The pair $(\mathbf{\Gamma}, D)$ is* implementable*, with complexity $Ng(T)$, if there exists an algorithm, running in time $\mathcal{O}(g(T))$, which takes $j \le N$ and $\alpha_j \in \operatorname{supp}(D)$ as input, and returns a weighted dataset, denoted as $S_j(\alpha_j)$, as output. The dataset $S_j(\alpha_j)$ must have a mass $\mathcal{O}(g(T))$ and* implement *the perturbation α_j in the following sense:*

$$\text{for all } x, x' \in \mathcal{X}\text{:} \quad \alpha_j(\Gamma_{xj} - \Gamma_{x'j}) = \sum_{(w,y) \in S_j(\alpha_j)} w\big(f(x, y) - f(x', y)\big).$$

This definition implies that the perturbation in the total utility described by the random vector $\boldsymbol{\alpha}$ can be imitated by appending a "fake history", consisting of the union of sets $S_j(\alpha_j)$, to the actual history of the game. While the actual value of a payoff for a given action x on such dataset might not agree with the perturbed objective, the difference of payoffs for any two actions x and x' agrees with the difference on the perturbed objective, and this suffices to find x_t.

One simple but powerful example of implementability is when the perturbations α_j are non-negative and each column of $\mathbf{\Gamma}$ corresponds to some adversary action y_j:

Example 3.10. *Consider any D with $\operatorname{supp}(D) \subseteq \mathbb{R}^+$. Suppose that for each $j \le N$, there exists some $y_j \in \mathcal{Y}$ such that $\Gamma_{xj} = f(x, y_j)$. Then $(\mathbf{\Gamma}, D)$ is implementable using the map $S_j(\alpha_j) = \{(\alpha_j, y_j)\}$.*

Algorithm 2: Oracle-Based Generalized FTPL

Input: positive $|\mathcal{X}| \times N$, an implementable $(\mathbf{\Gamma}, D)$, an offline oracle OPT
Draw $\alpha_j \sim D$ for $j = 1, \ldots, N$.
for $t = 1, \ldots, T$ **do**
  Set $S = \{(1, y_1), \ldots, (1, y_{t-1})\} \cup \bigcup_{j \le N} S_j(\alpha_j)$.
  Set $x_t = \text{OPT}(S)$.
  Observe y_t and receive payoff $f(x_t, y_t)$.
end for

Using implementable $(\mathbf{\Gamma}, D)$ gives rise to an oracle-efficient variant of the Generalized FTPL, provided in Algorithm 2. Theorem 3.11, with a proof in Appendix B, shows that the output of this algorithm is equivalent to the output of Generalized FTPL, and therefore the same regret guarantees hold. This result can guide the choice of $\mathbf{\Gamma}$ and D to obtain oracle-efficient no-regret algorithms.

Theorem 3.11. *If $(\mathbf{\Gamma}, D)$ is implementable with complexity $g(N, T)$ then Algorithm 2 is an oracle-efficient implementation of Algorithm 1 with per-round complexity $\mathcal{O}(T + g(N, T))$.*

4 Online Auction Design

In this section, we apply the general techniques developed in Section 3 to obtain efficient oracle-based no-regret algorithms for several common auction classes.

Consider a mechanism design setting where a seller wants to decide the allocation of resources to a set of n bidders. An allocation of a bidder is a subset in $\{0,1\}^k$ from some number of elements k and the seller has some feasibility constraints on the allocations across bidders. Each bidder $i \in [n]$ has a combinatorial valuation function $v_i \in \mathcal{V}$, where $\mathcal{V} \subseteq (\{0,1\}^k \to [0,1])$. A special case of the setting is that of multi-item auctions for k heterogeneous items, where each element is an item and the feasibility constraint simply states that no item is allocated to more than one bidder. Another special case is that of service based environments where each bidder's allocation is in $\{0,1\}$ and the seller has some constraints on which bidders to serve, i.e. assign an allocation of 1. Auction a takes as input the *bid profile* $\mathbf{b}$ consisting of reported values for each bidder, and returns both the allocation of bidder i, $\mathbf{q}_i(\mathbf{b}) \in \{0,1\}^k$ and the price $p_i(\mathbf{b})$ that this bidder is charged; we allow sets $\mathbf{q}_i(\mathbf{b})$ to overlap across bidders i.

We consider bidders with quasilinear utilities: the utility of bidder i is $v_i(\mathbf{q}_i(\mathbf{b})) - p_i(\mathbf{b})$. For an auction a with price function $\mathbf{p}(\cdot)$, we denote by $\text{Rev}(a, \mathbf{b})$ the *revenue of the auction* for bid profile $\mathbf{b}$, i.e., $\text{Rev}(a, \mathbf{b}) = \sum_{i \in [n]} p_i(\mathbf{b})$.

For single-parameter service based environments (a special case of which are single-item auctions), we slightly simplify notation and use $v_i \in [0,1]$ to denote the value of bidder i for being served.

In this work, we only consider *truthful auctions*, where each bidder i maximizes his utility by reporting his true valuation v_i, irrespective of what other bidders report. Since we limit attention to truthful auctions, we make the assumption that $\mathbf{b} = \mathbf{v}$ and refer to $\mathbf{v}$ as the *bid profile* throughout the rest of this section.

Fixing a class of (truthful) auctions $\mathcal{A}$ and a set of possible valuations $\mathcal{V}$, we consider the problem in which on each round $t = 1, \ldots, T$, a learner chooses an auction $a_t \in \mathcal{A}$ while an adversary chooses a bid profile $\mathbf{v}_t \in \mathcal{V}^n$. The learner then observes $\mathbf{v}_t$ and receives revenue $\text{Rev}(a_t, \mathbf{v}_t)$. The goal of the learner is to obtain low expected regret with respect to the best auction from $\mathcal{A}$ in hindsight. That is, we would like to guarantee that

$$\text{REGRET} = \mathbb{E}\left[\max_{a \in \mathcal{A}} \sum_{t=1}^{T} \text{Rev}(a, \mathbf{v}_t) - \sum_{t=1}^{T} \text{Rev}(a_t, \mathbf{v}_t)\right] \le o(T)\text{poly}(n, k).$$

We require our online algorithm to be oracle-efficient, assuming access to an offline optimization oracle that takes as an input a weighted set of bid profiles, $S = \{(w_\ell, \mathbf{v}_\ell)\}_{\ell \in \mathcal{L}}$ and returns the auction that maximizes the revenue on S, i.e., $\arg\max_{a \in \mathcal{A}} \sum_{(w,\mathbf{v}) \in S} w\mathrm{Rev}(a, \mathbf{v})$.

Using the language of oracle-based online learning developed in Section 3, the learner's action corresponds to the choice of auction, the adversary's action corresponds to the choice of bid profile, the payoff of the learner corresponds to the revenue generated by the auction, and we assume access to an offline optimization oracle OPT. These correspondences are summarized in the following table.

Auction Setting	**Oracle-Based Learning Equivalent**
Auctions $a_t \in \mathcal{A}$	Learner actions $x_t \in \mathcal{X}$
Bid/valuation profiles $\mathbf{v}_t \in \mathcal{V}^n$	Adversary actions $y_t \in \mathcal{Y}$
Revenue function Rev	Payoff function f

For several of the auction classes we consider, such as multi-item or multi-unit auctions, the revenue of an auction on a bid profile is in range $[0, R]$ for $R > 1$. In order to use the results of Section 3, we implicitly re-scale all the revenue functions and $\mathbf{\Gamma}$, by dividing them by R, before applying Theorem 3.5. Note that, while the admissibility condition also scales, the scaling of $\gamma = \max_{a,j} \mathbf{\Gamma}_{a,j}$ keeps the regret of the normalized problem at $O(N\sqrt{T\kappa\gamma/\delta})$, according to Theorem 3.5. We then scale up to get a regret bound that is R times the regret for the normalized problem, i.e., $O(RN\sqrt{T\kappa\gamma/\delta})$. This re-scaling increases the runtime by a factor of R, as when both the revenues and matrix $\mathbf{\Gamma}$ are scaled, the implementability map $S_j(\alpha_j)$ remains the same, and the support of the distribution (corresponding to $1/\eta$ in Theorem 3.5) increases by a factor of R.

We now derive results for VCG auctions with bidder-specific reserves, envy-free item-pricing auctions, and level auctions. We defer the definition of each auction class to its respective subsection.

4.1 VCG with Bidder-Specific Reserves

In this section, we consider a standard class of auctions, VCG auctions with bidder-specific reserve prices, which we define more formally below and denote by $\mathcal{I}$. These auctions are known to approximately maximize the revenue when the bidder valuations are drawn independently (but not necessarily identically) from some distribution [21]. Recently, Roughgarden and Wang [33] considered this class $\mathcal{I}$ in an online learning framework. They provided a computationally efficient algorithm whose total revenue is at least $1/2$ of the best revenue among auctions in $\mathcal{I}$, plus a term that is $o(T)$. We apply the techniques from Section 3 to generate an oracle-efficient online algorithm with low *additive* regret to the optimal auction in the class $\mathcal{I}$, without any loss in multiplicative factors.

We go beyond single-item auctions and consider a general *single-parameter setting*. In these environments, each bidder has one piece of private valuation for receiving a *service*, i.e., being included in the set of winning bidders. We allow for some combinations of bidders to be *served* simultaneously, and let $\mathcal{S} \subseteq 2^n$ be the family of feasible sets, i.e., sets of bidders that can be served simultaneously. We assume that it is possible for any bidder to be the sole bidder served, i.e., that $\{i\} \in \mathcal{S}$ for all i, and that it is possible that no bidder is served, i.e., $\emptyset \in \mathcal{S}$.[4] Examples of such environments include single-item single-unit auctions (for which $\mathcal{S}$ contains only singletons and the empty set), single-item s-unit auctions (for which $\mathcal{S}$ contains any subset of size at most s), and combinatorial auctions with single-minded bidders. In the last case, we begin with some set of original items, define the service as receiving the desired bundle of items, and let $\mathcal{S}$ contain any subset of bidders seeking disjoint sets of original items.

We consider the class of VCG auctions with bidder-specific reserves. In a basic VCG auction, an allocation $\mathbf{q}^*$ is chosen to maximize social welfare, that is, maximize $\sum_{i=1}^n v_i(q_i^*)$. Each bidder who is served is then charged the externality he imposes on others, $p_i(\mathbf{v}) = \max_{\mathbf{q}} \sum_{i' \neq i} v_{i'}(q_{i'}) - \sum_{i' \neq i} v_{i'}(q_{i'}^*)$, which can

[4] A more common and stronger assumption used in previous work [21, 33] is that $\mathcal{S}$ is a downward closed matroid.

be shown to equal the minimum bid at which he would be served. Such auctions are known to be truthful. The most common example is the second-price auction for the single-item single-unit case in which the bidder with the highest bid receives the item and pays the second highest bid. VCG auctions with reserves, which maintain the property of truthfulness, are defined as follows.

Definition 4.1 (VCG auctions with bidder-specific reserves)**.** *A VCG auction with bidder-specific reserves is specified by a vector* $\mathbf{r}$ *of reserve prices for each bidder. As a first step, all bidders whose bids are below their reserves (that is, bidders* i *for which* $v_i < r_i$*) are removed from the auction. If no bidders remain, the item is not allocated. Otherwise, the basic VCG auction is run on the remaining bidders to determine the allocation. Each bidder who is served is charged the larger of his reserve and his VCG payment.*

Fixing the set $\mathcal{S}$ of feasible allocations, we denote by $\mathcal{I}$ the class of all VCG auctions with bidder-specific reserves. With a slight abuse of notation we write $\mathbf{r} \in \mathcal{I}$ to denote the auction with reserve prices $\mathbf{r}$. To apply the results from Section 3, which require a finite action set for the learner, we limit attention to the finite set of auctions $\mathcal{I}_m \subseteq \mathcal{I}$ consisting of those auctions in which the reserve price for each bidder is an integer multiple of $1/m$. As we show below, the best auction in this class yields almost as high revenue as the best auction in $\mathcal{I}$ for common choices of $\mathcal{S}$.

We show how to design a matrix $\mathbf{\Gamma}$ and distribution D for this problem that are admissible and implementable. As a warmup, suppose we use the $|\mathcal{I}_m| \times n$ matrix $\mathbf{\Gamma}$ with $\Gamma_{\mathbf{r}i} = \mathrm{Rev}(\mathbf{r}, \mathbf{e}_i)$ for all $\mathbf{r} \in \mathcal{I}_m$ and $i \in [n]$. That is, the i^{th} column of $\mathbf{\Gamma}$ corresponds to the revenue of each auction on a bid profile in which bidder i has valuation 1 and all others have valuation 0. Let D be defined as in Theorem 3.5. By definition, $(\mathbf{\Gamma}, D)$ is implementable. Moreover, $\mathrm{Rev}(\mathbf{r}, \mathbf{e}_i) = r_i$ so any two rows of $\mathbf{\Gamma}$ are indeed different and $\mathbf{\Gamma}$ is $(m, 1/m)$-admissible. By Theorem 3.5, there is an oracle-efficient implementation of the Generalized FTPL with regret that is polynomial in m. In what follows, we design another translation matrix more carefully, and obtain the regret that is polynomial $\log(m)$.

Construction of $\mathbf{\Gamma}$: Let $\mathbf{\Gamma}^{\mathrm{VCG}}$ be an $|\mathcal{I}_m| \times (n\lceil \log m \rceil)$ binary matrix, where the i^{th} collection of $\lceil \log m \rceil$ columns correspond to the binary encoding of each auction's reserve price for bidder i. More formally, for any $i \le n$ and a bit position $\beta \le \lceil \log m \rceil$, $\mathbf{\Gamma}^{\mathrm{VCG}}_{\mathbf{r}j}$ is the β^{th} bit of (the integer) mr_i, where $j = (i-1)\lceil \log m \rceil + \beta$. The next lemma shows that this choice of $\mathbf{\Gamma}^{\mathrm{VCG}}$ is admissible and that $(\mathbf{\Gamma}^{\mathrm{VCG}}, D)$ is implementable for any D over $\mathbb{R}$.

Lemma 4.2. $\mathbf{\Gamma}^{\mathrm{VCG}}$ *is* $(2,1)$*-admissible. If* D *is supported on a subset of* $[-\alpha_{\max}, \alpha_{\max}]$*, then* $(\mathbf{\Gamma}^{\mathrm{VCG}}, D)$ *is implementable with complexity* $\mathrm{poly}(n, m, \alpha_{\max})$.

Proof. In the interest of readability, we drop the superscript and write $\mathbf{\Gamma}$ for $\mathbf{\Gamma}^{\mathrm{VCG}}$ in this proof.

For any $\mathbf{r}$, row $\Gamma_{\mathbf{r}}$ corresponds to the binary encoding of $r_1, \ldots, r_n$. Therefore, for any two different auctions $\mathbf{r} \neq \mathbf{r}'$, $\mathbf{\Gamma}_{\mathbf{r}} \neq \mathbf{\Gamma}_{\mathbf{r}'}$. Since $\mathbf{\Gamma}$ is a binary matrix, this implies that $\mathbf{\Gamma}$ is $(2,1)$-admissible.

Next, we prove that $(\mathbf{\Gamma}, D)$ is implementable. Pick $i \le n$ and $\beta \le \lceil \log m \rceil$, and the associated column index j. We will construct the set $S_j(\alpha_j)$ which implements the perturbation α_j. The set $S_j(\alpha_j)$ includes exactly the m profiles in which only the bidder i has non-zero valuation, denoted as $\mathbf{v}_h := (h/m)\mathbf{e}_i$ for $h \le m$. To determine their weights w_h, we use the definition of implementability. In particular, the weights must satisfy:

$$\forall\, \mathbf{r}, \mathbf{r}' \in \mathcal{I}_m, \qquad \alpha_j \left(\Gamma_{\mathbf{r}j} - \Gamma_{\mathbf{r}'j}\right) = \sum_{h \le m} w_h \Big(\mathrm{Rev}(\mathbf{r}, \mathbf{v}_h) - \mathrm{Rev}(\mathbf{r}', \mathbf{v}_h) \Big).$$

In the above equation, $\Gamma_{\mathbf{r}j}$ and $\Gamma_{\mathbf{r}'j}$ encode the β^{th} bit of r_i and r'_i, respectively, so the left-hand side is independent of the reserve prices for bidders $i' \neq i$. Moreover, $\mathrm{Rev}(\mathbf{r}, \mathbf{v}_h) = r_i \mathbf{1}_{(h \ge mr_i)}$, so the right-hand side of the above equation is also independent of the reserve prices for bidders $i' \neq i$. Let k_β be the β^{th}

bit of integer k. That is, $\Gamma_{\mathbf{r}j} = (mr_i)_\beta$. Substituting $k = mr_i$ and $k' = mr'_i$, the above equation can be reformulated as

$$\forall k, k' \in \{1, \dots, m\}, \quad \alpha_j \left(k_\beta - k'_\beta\right) = \sum_{h \le m} w_h \left(\frac{k}{m}\mathbf{1}_{(h \ge k)} - \frac{k'}{m}\mathbf{1}_{(h \ge k')}\right). \tag{5}$$

We next recursively derive the weights w_h, and show that they are non-negative and satisfy Equation (5). To begin, let

$$w_m = \max\Big\{0,\ \max_k \big[\alpha_j m\big(k_\beta - (k-1)_\beta\big)\big]\Big\},$$

and for all $k = m, m-1, \dots, 2$, define

$$w_{k-1} = \frac{1}{k-1}\left(\sum_{h=k}^{m} w_h - \alpha_j m\big(k_\beta - (k-1)_\beta\big)\right).$$

By definition, $w_m \ge 0$. Assume that for all $h \ge k$, $w_h \ge 0$. Then

$$w_{k-1} \ge \frac{1}{k-1}\Big(w_m - \alpha_j m\big(k_\beta - (k-1)_\beta\big)\Big) \ge 0.$$

Therefore all weights are non-negative. Furthermore, by rearranging the definition of w_{k-1}, we have

$$\begin{aligned}\alpha_j\big(k_\beta - (k-1)_\beta\big) &= \frac{1}{m}\left(\sum_{h=k}^{m} w_h - (k-1)w_{k-1}\right) = \frac{1}{m}\left(k\sum_{h=k}^{m} w_h - (k-1)\sum_{h=k-1}^{m} w_h\right) \\ &= \sum_{h \le m} w_h \left(\frac{k}{m}\mathbf{1}_{(h \ge k)} - \frac{k-1}{m}\mathbf{1}_{(h \ge k-1)}\right).\end{aligned}$$

Equation (5) is proved for a particular pair $k > k'$ by summing the above expression for $\alpha_j\big(s_\beta - (s-1)_\beta\big)$ over all $s \in (k', k]$ and canceling telescoping terms, and if $k = k'$, the statement holds regardless of the weights chosen.

This shows that $(\mathbf{\Gamma}, D)$ is implementable. It remains to argue the complexity. The running time of the construction of the weights is clearly $\mathrm{poly}(n, m)$ and the cardinality of $S_j(\alpha_j)$ is m. The remaining piece is bounding the sum of weights as $\mathrm{poly}(n, m, \alpha_{\max})$, which we defer to Appendix D.2. □

The next theorem follows immediately from Lemma 4.2, Theorem 3.5, and the fact that the maximum revenue is at most R.

Theorem 4.3. *Consider the online auction design problem for the class of VCG auctions with bidder-specific reserves, $\mathcal{I}_m$. Let $R = \max_{\mathbf{r},\mathbf{v}} \mathrm{Rev}(\mathbf{r}, \mathbf{v})$ and let D be the uniform distribution as described in Theorem 3.5. Then, the Oracle-Based Generalized FTPL algorithm with $\mathbf{\Gamma}^{\mathrm{VCG}}$ and D is oracle-efficient with complexity $\mathrm{poly}(n, m, T)$ and has regret*

$$\mathbb{E}\left[\max_{\mathbf{r} \in \mathcal{I}_m} \sum_{t=1}^{T} \mathrm{Rev}(\mathbf{r}, \mathbf{v}_t) - \sum_{t=1}^{T} \mathrm{Rev}(\mathbf{r}_t, \mathbf{v}_t)\right] \le O(n \log(m) R\sqrt{T}).$$

Note that in general, R is bounded by the number of bidders that can be served simultaneously, which is at most n.

Now let us return to the infinite class $\mathcal{I}$ of all VCG auctions with reserve prices $r_i \in [0, 1]$. We show $\mathcal{I}_m$ is a finite "cover" for this class when the family of feasible sets $\mathcal{S}$ is the set of all subsets of size at most s, corresponding to single-item single-unit auctions (when $s = 1$) or more general single-item s-unit auctions.

Consider any vector of reserves $\mathbf{r} \in \mathcal{I}$ and let $\mathbf{r}' \in \mathcal{I}_m$ be the vector obtained by rounding each reserve price down to the nearest multiple of $1/m$. If $v_i > r_i$, then $v_i > r'_i$, so any bidder who would have been included in the basic VCG auction using reserves $\mathbf{r}$ is still included with $\mathbf{r}'$. This can only increase the number of bidders who are serviced and therefore pay a charge. It is not hard to see, and we prove in Appendix D.1, that each bidder's payment can decrease by at most $1/m$. That is,

$$\max_{\mathbf{r}\in\mathcal{I}} \sum_{t=1}^{T} \mathrm{Rev}(\mathbf{r}, \mathbf{v}_t) - \max_{\mathbf{r}\in\mathcal{I}_m} \sum_{t=1}^{T} \mathrm{Rev}(\mathbf{r}, \mathbf{v}_t) \leq \frac{Ts}{m}. \tag{6}$$

Setting $m = \sqrt{T}$ and using Theorem 4.3, we obtain the following result for the class of auctions $\mathcal{I}$.

Theorem 4.4. *Consider the online auction design problem for the class of VCG auctions with bidder-specific reserves, $\mathcal{I}$, in s-unit auctions. Let D be the uniform distribution as described in Theorem 3.5. Then, the Oracle-Based Generalized FTPL algorithm with $\mathbf{\Gamma}^{\mathrm{VCG}}$ and D is oracle-efficient with complexity* $\mathrm{poly}(n, T)$ *and has regret efficient oracle-based online algorithm with regret*

$$\mathbb{E}\left[\max_{\mathbf{r}\in\mathcal{I}} \sum_{t=1}^{T} \mathrm{Rev}(\mathbf{r}, \mathbf{v}_t) - \sum_{t=1}^{T} \mathrm{Rev}(\mathbf{r}_t, \mathbf{v}_t)\right] \leq O(ns\log(T)\sqrt{T}).$$

4.2 Envy-free Item Pricing

In this section, we consider envy-free item pricing [17] in an environment with k heterogeneous items with a supply of $s_\ell \geq 0$ units for each item $\ell \leq k$.

Definition 4.5 (Envy-free Item-Pricing Auction)**.** *An* envy-free item-pricing auction *for k heterogeneous items, given supply s_ℓ for $\ell = 1, \ldots, k$ is defined by a vector of prices $\mathbf{a}$, where a_ℓ is the price of item ℓ. The mechanism considers bidders $i = 1, \ldots, n$ in order and allocates to bidder i the bundle $\mathbf{q}_i \in \{0, 1\}^k$ that maximizes $v_i(\mathbf{q}_i) - \mathbf{a}\cdot\mathbf{q}_i$, among all feasible bundles, i.e., bundles that can be composed from the remaining supplies. Bidder i is then charged the price $\mathbf{a}\cdot\mathbf{q}_i$.*

Examples of such environments include *unit-demand* bidders in *unlimited supply* setting and *single-minded* bidders in *unlimited supply* setting, such as *hypergraph pricing*, where bidders seek hyper-edges in a hypergraph, and its variant *the highway problem*, where bidders seek hyperedges between sets of contiguous vertices [2, 17]. We will describe some of these problems in more detail later on.

We represent by $\mathcal{P}_m$ the class of all such envy-free item pricing auctions where a_ℓ is a multiple of $1/m$ for all ℓ. Next, we discuss the construction of an implementable and admissible pair $(\mathbf{\Gamma}, D)$. Consider a bid profile where one bidder has value v for bundle $\mathbf{e}_\ell$ and all other bidders have value 0 for all bundles. The revenue of auction $\mathbf{a}$ on such a bid profile is $a_\ell \mathbf{1}_{(v\geq a_\ell)}$. Note the similarity to the case of VCG auctions with bidder-specific reserve prices $\mathbf{r}$, where bid profiles with a single non-zero valuation v_i and revenue $r_i\mathbf{1}_{(v_i\geq r_i)}$ were used to create an implementable construction for $\mathbf{\Gamma}$. We show that a similar construction works for $\mathcal{P}_m$.

Construction of $\mathbf{\Gamma}$: Let $\mathbf{\Gamma}^{\mathrm{IP}}$ be a $|\mathcal{P}_m| \times (k\lceil\log m\rceil)$ binary matrix, where the ℓ^{th} collection of $\lceil\log m\rceil$ columns correspond to the binary encoding of auction's price for item ℓ. More formally, for any $\ell \leq k$ and $\beta \leq \lceil\log m\rceil$, $\Gamma^{\mathrm{IP}}_{\mathbf{a}j}$ is the β^{th} bit of (the integer) ma_ℓ, where $j = (\ell-1)\lceil\log m\rceil + \beta$. Next, we show that $\mathbf{\Gamma}^{\mathrm{IP}}$ is admissible and $(\mathbf{\Gamma}^{\mathrm{IP}}, D)$ is implementable for any D with a bounded support. The proof of the following lemma is equivalent to that of Lemma 4.2 and appears in Appendix D.3 for completeness.

Lemma 4.6. *$\mathbf{\Gamma}^{\mathrm{IP}}$ is $(2, 1)$-admissible. If D is supported on a subset of $[-\alpha_{\max}, \alpha_{\max}]$, then $(\mathbf{\Gamma}^{\mathrm{IP}}, D)$ is implementable with complexity* $\mathrm{poly}(k, m, \alpha_{\max})$.

Our main theorem follows immediately from Lemma 4.6, Theorem 3.5, and the fact that the revenue of the mechanism at every step is at most R. In general, R is at most n.

Theorem 4.7. *Consider the online auction design problem for the class of envy-free item pricing auctions, $\mathcal{P}_m$. Let $R = \max_{\mathbf{a},\mathbf{v}} \mathrm{Rev}(\mathbf{a}, \mathbf{v})$ and let D be the uniform distribution as described in Theorem 3.5. Then, the Oracle-Based Generalized FTPL algorithm with $\mathbf{\Gamma}^{\mathrm{IP}}$ and D is oracle-efficient with complexity $\mathrm{poly}(n, k, m, T)$ and has regret*

$$\mathbb{E}\left[\max_{\mathbf{a}\in\mathcal{P}_m} \sum_{t=1}^{T} \mathrm{Rev}(\mathbf{a}, \mathbf{v}_t) - \sum_{t=1}^{T} \mathrm{Rev}(\mathbf{a}_t, \mathbf{v}_t)\right] \leq O\left(kR\log(m)\sqrt{T}\right).$$

Consider the class of all envy-free item-pricing auctions where $a_\ell \in [0, 1]$ is a real number and denote this class by $\mathcal{P}$. We show that $\mathcal{P}_m$ is a discrete "cover" for $\mathcal{P}$ when there is an *unlimited supply* of all items ($s_\ell = \infty$ for all ℓ) and the bidders have *single-minded* or *unit-demand* valuations.

In the single-minded setting, each bidder i is interested in one particular bundle of items $\hat{\mathbf{q}}_i$. That is, $v_i(\mathbf{q}_i) = 0$ for all $\mathbf{q}_i \neq \hat{\mathbf{q}}_i$. Consider any vector of prices $\mathbf{a} \in \mathcal{P}$ and let $\mathbf{a}' \in \mathcal{P}_m$ be the vector obtained by rounding each price down to the nearest multiple of $1/m$. Since the price of every bundle is reduced, any bidder i who received bundle $\hat{\mathbf{q}}_i$ in auction $\mathbf{a}$, also receives $\hat{\mathbf{q}}_i$ in auction $\mathbf{a}'$. So, the revenue of the mechanism reduces by at most nk/m. That is,

$$\max_{\mathbf{a}\in\mathcal{P}} \sum_{t=1}^{T} \mathrm{Rev}(\mathbf{a}, \mathbf{v}_t) - \max_{\mathbf{a}\in\mathcal{P}_m} \sum_{t=1}^{T} \mathrm{Rev}(\mathbf{a}, \mathbf{v}_t) \leq \frac{Tnk}{m}.$$

The following result is an immediate consequence of the above discretization for the choice of $m = \sqrt{T}$, the fact that $R \leq n$ for single-minded bidders, and Theorem 4.7.

Theorem 4.8. *Consider the online auction design problem for the class of envy-free item pricing auctions, $\mathcal{P}$, with single-minded bidders. Let D be the uniform distribution as described in Theorem 3.5. Then, the Oracle-Based Generalized FTPL algorithm with $\mathbf{\Gamma}^{\mathrm{IP}}$ and D is oracle-efficient with complexity $\mathrm{poly}(k, n, T)$ and has regret*

$$\mathbb{E}\left[\max_{\mathbf{a}\in\mathcal{P}} \sum_{t=1}^{T} \mathrm{Rev}(\mathbf{a}, \mathbf{v}_t) - \sum_{t=1}^{T} \mathrm{Rev}(\mathbf{a}_t, \mathbf{v}_t)\right] \leq O\left(nk\log(T)\sqrt{T}\right).$$

In the unit-demand setting, each bidder i has valuation $v_i(\mathbf{e}_\ell)$ for item ℓ, and wishes to purchase *at most one item*, i.e., item $\arg\max_\ell (v_i(\mathbf{e}_\ell) - a_\ell)$. We show that for any $\mathbf{a} \in \mathcal{P}$ there is $\mathbf{a}' \in \mathcal{P}_m$ such that for any valuation profile $\mathbf{v}$, $\mathrm{Rev}(\mathbf{a}, \mathbf{v}) - \mathrm{Rev}(\mathbf{a}', \mathbf{v}) \leq O(nk/m)$. At a high level, we choose $\mathbf{a}'$ as discounted prices such that items with higher prices are discounted at higher rates. It is not hard to see that under this condition, no bidder purchases a less expensive item in auction $\mathbf{a}'$. So, the loss in the revenue of the auction is bounded by the discount on the items. Using this intuition, it is sufficient to find $\mathbf{a}' \in \mathcal{P}_m$ such that $a_\ell \geq a'_\ell \geq a_\ell - O(1/m)$ for all $\ell \leq k$, and $a_\ell - a'_\ell > a_{\ell'} - a'_{\ell'}$ when $a_\ell > a_{\ell'}$. In the next lemma, we provide one such mapping between $\mathbf{a}$ and $\mathbf{a}' \in \mathcal{P}_m$ that has an additive approximation guarantee. Hartline and Koltun [20] also provided a discretization of $\mathcal{P}$ that has multiplicative approximation guarantee, but using a discretized set different from $\mathcal{P}_m$.

Lemma 4.9. *For any $\mathbf{a} \in \mathcal{P}$ there is $\mathbf{a}' \in \mathcal{P}_m$, such that for any valuation profile $\mathbf{v}$, $\mathrm{Rev}(\mathbf{a}, \mathbf{v}) - \mathrm{Rev}(\mathbf{a}', \mathbf{v}) \leq nk/m$.*

Proof. For ease of exposition, let $\epsilon = 1/m$. Without loss of generality, assume $a_1 \leq a_2 \leq \cdots \leq a_k$. We next derive the desired $\mathbf{a}'$. To begin, let a'_1 be the largest multiple of ϵ less than or equal to a_1. For $\ell = 2, \ldots, k$, let a'_ℓ be the largest multiple of ϵ less than or equal to a_ℓ such that $a_\ell - a'_\ell \geq a_{\ell-1} - a'_{\ell-1}$. Note that a'_ℓ is well defined, since we can begin by considering $a'_\ell = a'_{\ell-1}$ and then increase by ϵ until the condition is violated. This construction means that pricing of items with a larger ℓ is more attractive in $\mathbf{a}'$ than it was

in $\mathbf{a}$. Thus, no bidder that prefers an item ℓ under $\mathbf{a}$, would prefer any item $\ell' < \ell$ under $\mathbf{a}'$. Therefore, the revenue obtained from the bidder i who prefers ℓ under $\mathbf{a}$ is at least a'_ℓ under $\mathbf{a}'$, which implies the bound

$$\mathrm{Rev}(\mathbf{a}, \mathbf{v}) - \mathrm{Rev}(\mathbf{a}', \mathbf{v}) \leq n \max_{\ell \leq k} (a_\ell - a'_\ell).$$

To complete the proof, we argue by induction that $a_\ell - a'_\ell \leq \ell\epsilon$. This clearly holds for $\ell = 1$. For $\ell \geq 2$, the definition of a'_ℓ, in the absence of discretization to ϵ, would yield $a'_\ell = a_\ell - (a_{\ell-1} - a'_{\ell-1})$. With the discretization, we have

$$a'_\ell \geq a_\ell - (a_{\ell-1} - a'_{\ell-1}) - \epsilon \geq a_\ell - (\ell - 1)\epsilon - \epsilon = a_\ell - \ell\epsilon,$$

where we used the inductive hypothesis at $\ell - 1$. □

For a choice of $m = \sqrt{T}$, the fact that $R \leq n$, and Theorem 4.7 we have the following result for the class of auctions $\mathcal{P}$.

Theorem 4.10. *Consider the online auction design problem for the class of envy-free item pricing auctions, $\mathcal{P}$, with unit-demand bidders. Let D be the uniform distribution as described in Theorem 3.5. Then, the Oracle-Based Generalized FTPL algorithm with $\mathbf{\Gamma}^{\mathrm{IP}}$ and D is oracle-efficient with complexity* $\mathrm{poly}(n, k, T)$ *and has regret*

$$\mathbb{E}\left[\max_{\mathbf{a} \in \mathcal{P}} \sum_{t=1}^{T} \mathrm{Rev}(\mathbf{a}, \mathbf{v}_t) - \sum_{t=1}^{T} \mathrm{Rev}(\mathbf{a}_t, \mathbf{v}_t)\right] \leq O\left(nk \log(T)\sqrt{T}\right).$$

4.3 Level Auctions

We consider the class of *level auctions* introduced by Morgenstern and Roughgarden [28]. These auctions can achieve $(1-\epsilon)$-approximate revenue maximization, when the valuations of the bidders are drawn independently (but not necessarily identically) from a distribution [28], approximating the Myerson's optimal auction [29]. Using our tools, we derive oracle-efficient no-regret algorithms for this auction class.

The s-level auctions realize a single-item single-unit allocation as follows:

Definition 4.11. *An s-level auction $\boldsymbol{\theta}$ is defined by s thresholds for each bidder i, $0 \leq \theta^i_0 \leq \cdots \leq \theta^i_{s-1} \leq 1$. For any bid profile $\mathbf{v}$, we let $b^{\boldsymbol{\theta}}_i(v_i)$ denote the index b of the largest threshold $\theta^i_b \leq v_i$, or -1 if $v_i < \theta^i_0$. If $v_i < \theta^i_0$ for all i, the item is not allocated. Otherwise, the item goes to the bidder with the largest index $b^{\boldsymbol{\theta}}_i(v_i)$, breaking ties in favor of bidders with smaller i. The winner pays the price equal to the minimum bid that it could have submitted and still won the item.*

When it is clear from the context, we suppress the notation of $\boldsymbol{\theta}$ in $b^{\boldsymbol{\theta}}_i(v_i)$. We consider two classes of s-level auctions, $\mathcal{R}_{s,m}$ and $\mathcal{S}_{s,m}$, where $\mathcal{R}_{s,m}$ is the set of all auctions described by Definition 4.11 with thresholds that are multiples of $1/m$ and $\mathcal{S}_{s,m} \subseteq \mathcal{R}_{s,m}$ is the set of s-level auctions such that the thresholds for each i are distinct.

Let us first consider $\mathcal{S}_{s,m}$ and construct an admissible and implementable pair $(\mathbf{\Gamma}, D)$. Consider the bid profile in which the only non-zero bids are $v_n = \ell/m$, for some $0 \leq \ell \leq m$, and $v_i = 1$, for one bidder $i < n$. Note that bidder i wins the item in any such profile and pays θ^i_b for $b = \max\{0, b_n(v_n)\}$. Consider matrix $\mathbf{\Gamma}$ whose columns correspond to the revenue of auctions in $\mathcal{S}_{s,m}$ on such bid profiles, with an additional column corresponding to the bid profile $\mathbf{e}_n$. Clearly, $(\mathbf{\Gamma}, D)$ is implementable when $\mathrm{supp}(D) \subseteq \mathbb{R}^+$. As for admissibility, take $\boldsymbol{\theta} \in \mathcal{S}_{s,m}$ and the corresponding row $\mathbf{\Gamma}_{\boldsymbol{\theta}}$. Note that as $v_n = \ell/m$ increases for $\ell = 0, \ldots, m$, there is an increase in $b_n(\ell/m) = -1, 0, \ldots, s-1$, possibly skipping the initial -1. Once $b_n(v_n)$ reaches 0 and then increases to $s-1$, so does the revenue of $\boldsymbol{\theta}$, i.e., $\theta^i_{b_n(v_n)}$. Since any two consecutive thresholds of $\boldsymbol{\theta}$ are different, the thresholds of θ^i_b, for $b \geq 0$, and θ^n_b, for $b \geq 1$, can be reconstructed by

analyzing the revenue of the auction and the values of v_n at which the revenue changes. The remaining threshold θ_0^n is recovered by the bid profile $\mathbf{v} = \mathbf{e}_n$. This shows that any two rows of $\mathbf{\Gamma}$ are different and $\mathbf{\Gamma}$ is admissible. This reasoning is summarized in the following construction and the corresponding lemma, whose proof appears in Appendix D.4.

Construction of $\mathbf{\Gamma}$: For $i \in \{1, \dots, n-1\}$ and $\ell \in \{0, \dots, m\}$, let $\mathbf{v}^{i,\ell} = \mathbf{e}_i + (\ell/m)\mathbf{e}_n$. Let $V = \{\mathbf{v}^{i,\ell}\}_{i,\ell} \cup \{\mathbf{e}_n\}$. Let $\mathbf{\Gamma}^{\mathrm{SL}}$ be the matrix of size $|\mathcal{S}_{s,m}| \times |V|$ with entries indexed by $(\boldsymbol{\theta}, \mathbf{v}) \in \mathcal{S}_{s,m} \times V$, such that $\Gamma^{\mathrm{SL}}_{\boldsymbol{\theta},\mathbf{v}} = \mathrm{Rev}(\boldsymbol{\theta}, \mathbf{v})$.

Lemma 4.12. *Let D be a distribution supported on a subset of $[0, \alpha_{\max}]$. Then $(\mathbf{\Gamma}^{\mathrm{SL}}, D)$ is implementable with complexity $\mathrm{poly}(n, m, \alpha_{\max})$ and $\mathbf{\Gamma}^{\mathrm{SL}}$ is $(m+1,\ 1/m)$-admissible.*

Our next theorem is an immediate consequence of the implementability of $(\mathbf{\Gamma}^{\mathrm{SL}}, D)$ and Theorem 3.5.

Theorem 4.13. *Consider the online auction design problem for the class of s-level auctions with no repeated thresholds, $\mathcal{S}_{s,m}$. Let D be the uniform distribution as described in Theorem 3.5. Then, the Oracle-Based Generalized FTPL algorithm with $\mathbf{\Gamma}^{\mathrm{SL}}$ and D is oracle-efficient with complexity $\mathrm{poly}(n, m, T)$ and has regret*

$$\mathbb{E}\left[\max_{\boldsymbol{\theta} \in \mathcal{S}_{s,m}} \sum_{t=1}^{T} \mathrm{Rev}(\boldsymbol{\theta}, \mathbf{v}_t) - \sum_{t=1}^{T} \mathrm{Rev}(\boldsymbol{\theta}_t, \mathbf{v}_t)\right] \leq O(nm^2\sqrt{T}).$$

Next, we consider the class of auctions $\mathcal{R}_{s,m}$ and construct an admissible and implementable pair $(\mathbf{\Gamma}, D)$. As a warm-up, we demonstrate how the structure of $\mathcal{R}_{s,m}$ differs from $\mathcal{S}_{s,m}$ and argue that $\mathbf{\Gamma}^{\mathrm{SL}}$ is not admissible for $\mathcal{R}_{s,m}$. Take $\mathbf{v}$ such that $v_i = 1$ and v_n is some multiple of $1/m$. Similarly as before, as v_n increases there is also an increase in $b_n(v_n)$ and the revenue $\theta^i_{b_n(v_n)}$. However, the thresholds of bidder n are no longer distinct, so $b_n(v_n)$ can skip certain values, and as a result θ^i_b is not revealed for all b. In addition, when the thresholds of bidder i are not distinct, we have that even when $b_n(v_n)$ strictly increases, there might be no increase in the revenue, and as a result θ^n_b cannot be reconstructed.

To identify an admissible construction for $\mathbf{\Gamma}$, we first analyze the structure of auctions in $\mathcal{R}_{s,m}$. Let $B^i_b(\boldsymbol{\theta})$ denote the set of all bids v_i such that $b^{\boldsymbol{\theta}}_i(v_i) = b$; we refer to such sets as *buckets*, since they partition i's valuations. Given a bucket $B^i_b(\boldsymbol{\theta})$, we define its *competing bucket* as any $B^{i'}_{b'}(\boldsymbol{\theta})$ such that either $i' < i$ and $b' = b-1$ or $i' > i$ and $b' = b$. Note that the payment of a winning bidder i is θ^i_b if another bidder's bid falls into a competing bucket of $B^i_b(\boldsymbol{\theta})$. By definition, $B^i_b(\boldsymbol{\theta}) = \emptyset$ when $\theta^i_b = \theta^i_{b+1}$. The next lemma, whose proof appears in Appendix D.5, shows that we can restrict our attention to auctions for which empty buckets satisfy a strong structural property:

Lemma 4.14. *Let $\overline{\mathcal{R}}_{s,m} \subseteq \mathcal{R}_{s,m}$ be the set of auctions $\boldsymbol{\theta}$ with the following property: Every non-empty bucket $B^i_b(\boldsymbol{\theta})$ with $b \geq 1$ has a non-empty competing bucket. Then, for any $\boldsymbol{\theta} \in \mathcal{R}_{s,m}$, there is $\boldsymbol{\theta}' \in \overline{\mathcal{R}}_{s,m}$, such that*

$$\forall \mathbf{v} \in [0,1]^n, \quad \mathrm{Rev}(\boldsymbol{\theta}, \mathbf{v}) = \mathrm{Rev}(\boldsymbol{\theta}', \mathbf{v}).$$

Lemma 4.14 implies that any no-regret algorithm for the class $\overline{\mathcal{R}}_{s,m}$ is also a no-regret algorithm for the class $\mathcal{R}_{s,m}$. So, it suffices to analyze the class of auctions $\overline{\mathcal{R}}_{s,m}$.

Next, we construct an admissible and implementable pair $(\mathbf{\Gamma}, D)$ for $\overline{\mathcal{R}}_{s,m}$, by considering revenue of auctions on bid profiles in which the only non-zero bids are $v_i = 1$ and $v_{i'}$ that is a multiple of $1/m$. We show that for any two $\boldsymbol{\theta}, \boldsymbol{\theta}' \in \overline{\mathcal{R}}_{s,m}$, there is one such bid profile where the revenues of these auctions differ. At a high level, we prove this by taking $\theta^i_b \neq \theta'^i_b$ with a non-empty bucket $B^i_b(\boldsymbol{\theta})$ and showing that there is a bid $v_{i'}$ in one of the competing buckets that enforces two different prices $\theta^i_b \neq \theta^i_b$. This is further formalized in the construction below and Lemma 4.15, whose proof appears in Appendix D.6.

Construction of $\mathbf{\Gamma}$: Let V be the set of all bid profiles $\mathbf{v}^{i,i',\ell} = \mathbf{e}_i + (\ell/m)\mathbf{e}_{i'}$ for $i \neq i' \in \{1,\ldots,n\}$ and $\ell \in \{0,\ldots,m\}$. Let $\mathbf{\Gamma}^{\mathrm{RL}}$ be a $|\overline{\mathcal{R}}_{s,m}| \times |V|$ matrix with entries indexed by $(\boldsymbol{\theta}, \mathbf{v}) \in \overline{\mathcal{R}}_{s,m} \times V$, such that $\Gamma^{\mathrm{RL}}_{\boldsymbol{\theta},\mathbf{v}} = \mathrm{Rev}(\boldsymbol{\theta}, \mathbf{v})$.

Lemma 4.15. *Let D be a distribution supported on a subset of $[0, \alpha_{\max}]$. Then, $(\mathbf{\Gamma}^{\mathrm{RL}}, D)$ is implementable with the complexity $\mathrm{poly}(n, m, \alpha_{\max})$ and $\mathbf{\Gamma}^{\mathrm{RL}}$ is $(m+1,\ 1/m)$-admissible.*

The next theorem is a direct consequence of Lemmas 4.15 and 4.14.

Theorem 4.16. *Consider the online auction design problem for the class of s-level auctions with repeated thresholds, $\mathcal{R}_{s,m}$. Let D be the uniform distribution as described in Theorem 3.5. Then, the Oracle-Based Generalized FTPL algorithm with $\mathbf{\Gamma}^{\mathrm{RL}}$ and D is oracle-efficient with complexity $\mathrm{poly}(n, m, T)$ and has regret*

$$\mathbb{E}\left[\max_{\boldsymbol{\theta}\in\mathcal{R}_{s,m}} \sum_{t=1}^{T} \mathrm{Rev}(\boldsymbol{\theta}, \mathbf{v}_t) - \sum_{t=1}^{T} \mathrm{Rev}(\boldsymbol{\theta}_t, \mathbf{v}_t)\right] \leq O(n^2 m^2 \sqrt{T}).$$

5 Stochastic Adversaries and Stronger Benchmarks

So far our results apply to general adversaries, where the sequence of adversary actions are arbitrary and where we showed that the payoff of the learner is at least as large as the payoff of the best action in hindsight. Can we make stronger statements about the average payoff of a no-regret learning algorithm when we impose distributional assumptions on the sequence of the adversary?

We start with the easier setting where the actions of the learner are drawn i.i.d. and then we analyze the slightly more complex setting where the actions of the learner follow a fast-mixing Markov chain. For both settings we show that the average payoff of the learning algorithm is as large as the optimal expected payoff, in expectation over the i.i.d. distribution in the i.i.d. setting and over the stationary distribution in the Markovian setting.

When applied to the online optimal auction setting, combining these results with approximate optimality results of simple auctions such as s-level auctions or VCG with bidder-specific reserves, we get that the average revenue of our online learning algorithms compete with the revenue achieved by the unrestricted optimal auction for these distributional settings and not only with the best auction within the class over which our algorithms were learning.

5.1 Stochastic Adversaries

I.I.D. Adversary One extreme case is to assume that the adversary's action y_t at each iteration is drawn independently and identically from the same unknown distribution F. This leads to the i.i.d. learning setting. An easy application of the Chernoff-Hoeffding bound yields that for such a learning setting, the average payoff of a no-regret learner is at least as large as the best payoff one could achieve in expectation over the distribution F.

Lemma 5.1. *Suppose that $y_1, \ldots, y_T$ are i.i.d. draws from a distribution F. Then for any no-regret learning algorithm, with probability at least $1-\delta$,*

$$\frac{1}{T}\sum_{t=1}^{T} \mathbb{E}_{x_t}[f(x_t, y_t)] \geq \sup_{x\in\mathcal{X}} \mathbb{E}_{y\sim F}[f(x, y)] - \sqrt{\frac{\log(2/\delta)}{2T}} - \frac{\textsc{Regret}}{T}. \tag{7}$$

Proof. Let $x^* = \arg\sup_{x\in\mathcal{X}} \mathbb{E}_{y\sim F}[f(x, y)]$. By the definition of regret we have that for any $y_1, \ldots, y_T$,

$$\sum_{t=1}^{T} \mathbb{E}_{x_t}[f(x_t, y_t)] \geq \sup_{x\in\mathcal{X}} \sum_{t=1}^{T} f(x, y_t) - \textsc{Regret} \geq \sum_{t=1}^{T} f(x^*, y_t) - \textsc{Regret}. \tag{8}$$

Observe that the random variables $Z_t = f(x^*, y_t)$ are drawn i.i.d. with expected value $\mathbb{E}_{y\sim F}[f(x^*, y)]$ and are bounded in $[0, 1]$. Hence, by the Hoeffding bound, we get that with probability at least $1 - 2e^{-2T\epsilon^2}$:

$$\frac{1}{T}\sum_{t=1}^{T} f(x^*, y_t) \geq \mathbb{E}_{y\sim F}[f(x^*, y)] - \epsilon. \tag{9}$$

By setting $\epsilon = \sqrt{\frac{\log(2/\delta)}{2T}}$ and combining the two bounds we get the result. □

Markovian Adversary Suppose that the choice of the adversary y_t follows a stationary and reversible Markov process based on some transition matrix $P(y, y')$ with a stationary distribution F. Moreover, consider the case where the set $\mathcal{Y}$ is finite. For any Markov chain, the spectral gap γ is defined as the difference between the first and the second largest eigenvalue of the transition matrix P (the first eigenvalue always being 1). We will assume that this gap is bounded away from zero. The spectral gap of a Markov chain is strongly related with its mixing time. In this work we will specifically use the following result of Paulin [31], which is a Bernstein concentration inequality for sums of dependent random variables that are the outcome of a stationary Markov chain with spectral gap bounded away from zero. A Markov chain $y_1, \ldots, y_T$ is stationary if $y_1 \sim F$, where F is the stationary distribution. For simplicity, we focus on stationary chains, though similar results hold for non-stationary chains (see Paulin [31] and references therein).

Theorem 5.2 (Paulin [31], Theorem 3.8)**.** *Let $X_1, \ldots, X_z$ be a stationary and reversible Markov chain on a state space Ω, with stationary distribution F and spectral gap γ. Let $g : \Omega \to [0, 1]$ and $S_z = \frac{1}{z}\sum_{i=1}^{z} g(X_i)$. Then*

$$\Pr[|S_z - \mathbb{E}_{X\sim F}[g(X)]| > \epsilon] \leq 2\exp\left(-\frac{z\gamma\epsilon^2}{4 + 10\epsilon}\right). \tag{10}$$

Lemma 5.3. *Suppose that $y_1, \ldots, y_T$ are a stationary and reversible Markov chain based on a transition matrix P, with stationary distribution F and with spectral gap γ. Then for any no-regret learning algorithm, with probability at least $1 - \delta$:*

$$\frac{1}{T}\sum_{t=1}^{T} \mathbb{E}_{x_t}[f(x_t, y_t)] \geq \sup_{x\in\mathcal{X}} \mathbb{E}_{y\sim F}[f(x, y)] - \sqrt{\frac{14\log(2/\delta)}{\gamma T}} - \frac{\textsc{Regret}}{T}. \tag{11}$$

Proof. Let $x^* = \arg\sup_{x\in\mathcal{X}} \mathbb{E}_{y\sim F}[f(x, y)]$. By the definition of regret we have that for any sequence $y_1, \ldots, y_T$,

$$\sum_{t=1}^{T} \mathbb{E}_{x_t}[f(x_t, y_t)] \geq \sup_{x\in\mathcal{X}} \sum_{t=1}^{T} f(x, y_t) - \textsc{Regret} \geq \sum_{t=1}^{T} f(x^*, y_t) - \textsc{Regret}. \tag{12}$$

Since $y_1, \ldots, y_T$ are a Markov chain, by applying Theorem 5.2 to this chain and to the function $f(x^*, \cdot)$, we obtain that with probability at least $1 - 2\exp\left(-\frac{T\gamma\epsilon^2}{4+10\epsilon}\right)$,

$$\frac{1}{T}\sum_{t=1}^{T} f(x^*, y_t) \geq \mathbb{E}_{y\sim F}[f(x^*, y)] - \epsilon. \tag{13}$$

If we set $\epsilon = \sqrt{\frac{14\log(2/\delta)}{\gamma T}}$ then we have, either $\epsilon > 1$, in which case the inequality is trivial, since $f(x, y) \in [0, 1]$, or if $\epsilon \leq 1$, then $\epsilon = \sqrt{\frac{14\log(2/\delta)}{\gamma T}} \geq \sqrt{\frac{(4+10\epsilon)\log(2/\delta)}{\gamma T}}$, which implies that the inequality holds with probability $1 - \delta$. □

Example 5.4 (Sticky Markov Chain). *Consider a Markov chain where at every iteration y_t is equal to y_{t-1} with some probability $\rho \geq 1/2$ and with the remaining probability $(1-\rho)$ it is drawn independently from some fixed distribution F. It is easy to see that the stationary distribution of this chain is equal to F. We can bound the spectral gap of this Markov chain by Cheeger's bound. The Cheeger constant for a finite state, reversible Markov chain is defined as*

$$\Phi = \min_{S\subseteq\Omega:F(S)\leq 1/2} \frac{\sum_{x\in S}\sum_{y\in S^c} F(x)P(x,y)}{F(S)} = \min_{S\subseteq\Omega:F(S)\leq 1/2} \frac{\sum_{x\in S}\sum_{y\in S^c} F(x)(1-\rho)F(y)}{F(S)}$$
$$= \min_{S\subseteq\Omega:F(S)\leq 1/2} (1-\rho)\frac{F(S)\cdot F(S^c)}{F(S)} = \min_{S\subseteq\Omega:F(S)\leq 1/2} (1-\rho)F(S^c) \geq \frac{1-\rho}{2}$$

Moreover, by the Cheeger bound [8] we know that $\gamma \geq \frac{\Phi^2}{2} \geq \frac{(1-\rho)^2}{4}$. Thus we get that for such a sequence of adversary actions, with probability $1-\delta$,

$$\frac{1}{T}\sum_{t=1}^{T} \mathbb{E}_{x_t}[f(x_t,y_t)] \geq \sup_{x\in\mathcal{X}} \mathbb{E}_{y\sim F}[f(x,y)] - \frac{2}{1-\rho}\sqrt{\frac{14\log(2/\delta)}{T}} - \frac{\textsc{Regret}}{T} \tag{14}$$

5.2 Implications for Online Optimal Auction Design

Consider the online optimal auction design problem for a single item and n bidders. Suppose that the adversary who picks the valuation vectors $\mathbf{v}_1,\ldots,\mathbf{v}_T$, is Markovian and that the stationary distribution F of the chain is independent across players, i.e., the stationary distribution is a product distribution $F = F_1 \times \ldots \times F_n$.

Then we know that the optimal auction for this setting is what is known as Myerson's auction [29], which translates the players' values based on some monotone function ϕ, known as the ironed virtual value function and then allocates the item to the bidder with the highest virtual value, charging payments so that the mechanism is dominant strategy truthful.

A recent result of Morgenstern and Roughgarden [28] shows that level auctions approximate Myerson's auction in terms of revenue. In particular, if distributions F_i are bounded in $[1,H]$, then the class of s-level auctions with $s = \Omega(\frac{1}{\epsilon} + \log_{1+\epsilon} H)$, where the thresholds can be any real numbers, achieves expected revenue at least $(1-\epsilon)$ of the expected optimal revenue of Myerson's auction. We prove and use a similar result for s-level auctions that uses the discretized set $\mathcal{R}_{s,m}$. That is, when $s = O(1/\epsilon)$ and $m = O(1/\epsilon)$, we have

$$\max_{\boldsymbol{\theta}\in\mathcal{R}_{s,m}} \mathbb{E}_{\mathbf{v}\sim F}[\mathrm{Rev}(\boldsymbol{\theta},\mathbf{v})] \geq \textsc{Opt}(F) - \epsilon. \tag{15}$$

The proof of this equation appears in Appendix E. Combining the results in this section with the aforementioned results we get the following theorem:

Theorem 5.5 (Competing with Overall Optimal). *Consider the online optimal auction design problem for a single item among n bidders, where the sequence of valuation vectors $\mathbf{v}_1,\ldots,\mathbf{v}_T$, is Markovian, following a stationary and reversible Markov process, with a stationary distribution $F = F_1 \times \ldots \times F_n$ that is a product distribution across bidders, and with a spectral gap of $\gamma > 0$. Then the oracle-efficient online learning algorithm which optimizes over the set of s-level auctions $\mathcal{R}_{s,m}$ with $s = O(T^{1/6}n^{-2/3})$ and with a discretization of the threshold levels of size $m = O(T^{1/6}n^{-2/3})$, guarantees, with probability at least $1-\delta$, average expected revenue at least:*

$$\frac{1}{T}\sum_{t=1}^{T} \mathbb{E}_{\boldsymbol{\theta}_t}\left[\mathrm{Rev}(\boldsymbol{\theta}_t,\mathbf{v}_t)\right] \geq \textsc{Opt}(F) - \frac{1}{m} - \sqrt{\frac{14\log(2/\delta)}{\gamma T}} - O\left(\frac{m^2n^2}{\sqrt{T}}\right) \geq \textsc{Opt}(F) - O\left(\frac{n^{2/3}}{T^{1/6}}\right), \tag{16}$$

where $\textsc{Opt}(F)$ is the optimal revenue achievable by any dominant strategy truthful mechanism for valuation vector distribution F.

Example 5.6 (Valuation Shocks)**.** *Consider the setting where valuations of players in the beginning of time are drawn from some product distribution $F = F_1 \times \ldots \times F_n$. Then at every iteration with some probability ρ the valuations of all players remain the same as in the previous iteration, while with some probability $1-\rho$, there is a shock in the market and the valuations of the players are re-drawn from distribution F. As we analyzed in the previous section, the spectral gap of the Markov chain defined by this setting is at least $\gamma \geq \frac{(1-\rho)^2}{4}$. Thus we get a regret bound which depends inversely proportionally with the quantity $1-\rho$.*

Hence, our online learning algorithm achieves revenue that is close to the optimal revenue achievable by any dominant strategy truthful mechanism for the distribution F. However, it achieves this guarantee even if the valuations of the players are not drawn i.i.d. at every iteration and even if the learner does not know the distribution F or when the valuations of the players are going to be re-drawn or what the rate ρ of shocks in the markets is.

6 Contextual Online Learning (Side Information)

We now consider a generalization of the online optimal auction setting, where at every iteration the learner also observes contextual information. The context σ_t at each iteration comes from some abstract context space Σ. In this setting the learner wants to use such contextual information to improve his performance. Specifically, the goal of the learner is to compete with a set of policies Π, where a policy is a mapping from a context σ_t to an auction $\pi(\sigma_t)$.

Given a class of auctions $\mathcal{A}$, a set of possible valuations $\mathcal{V}$, a policy class Π and an unknown sequence of bid profiles and contexts $(\sigma_1, \mathbf{v}_1), \ldots, (\sigma_T, \mathbf{v}_T) \in \Sigma \times \mathcal{V}^n$, our goal is to provide an online algorithm that picks a policy $\pi_t \in \Pi$ at every time step and achieves a revenue that is close to the revenue of the best policy $\pi \in \Pi$ in hindsight. That is,

$$\mathbb{E}\left[\max_{\pi\in\Pi}\sum_{t=1}^{T}\mathrm{Rev}(\pi(\sigma_t),\mathbf{v}_t) - \sum_{t=1}^{T}\mathrm{Rev}(\pi_t(\sigma_t),\mathbf{v}_t)\right] \leq o(T).$$

6.1 Non-Contextual to Contextual Learning for Separable Policies

We will give a reduction from a non-contextual setting to a contextual one in the general learning setting rather than a specific analysis for the auction setting. Hence, we consider the learning setting of Section 3 in which the learner has an action space $\mathcal{X}$, the adversary has an action space $\mathcal{Y}$ and the payoff function of the learner is given by $f(x,y)$ for each $x \in \mathcal{X}$ and $y \in \mathcal{Y}$.

We will then consider the contextual version of this setting, where the learner also observes contextual information $\sigma \in \Sigma$, and picks a policy $\pi \in \Pi$ where Π is a subset of the mappings from a context σ to an action $x \in X$. The action space of the adversary is a context $\sigma \in \Sigma$ and an action $y \in \mathcal{Y}$. Then the payoff of the learner is given by:

$$f_c(\pi, (\sigma, y)) = f(\pi_t(\sigma), y_t) \tag{17}$$

while the regret of any algorithm is given by:

$$\textsc{Regret} = \mathbb{E}\left[\max_{\pi\in\Pi}\sum_{t=1}^{T} f(\pi(\sigma_t), y_t) - \sum_{t=1}^{T} f(\pi_t(\sigma_t), y_t)\right].$$

The offline oracle assumption now asks for an algorithm that takes as input a distribution over contexts and adversary actions and returns the best policy in the policy space for this distribution.

The contextual learning setting is another instance of the general learning setting, with learner action space $\mathcal{X}_c = \Pi \subseteq (\Sigma \to \mathcal{X})$, adversary action space $\mathcal{Y}_c = \Sigma \times \mathcal{Y}$ and payoff function $f_c(x_c, y_c)$ as defined above. Moreover, a *contextual offline oracle* is an offline oracle for this new learning setting.

If we are given an admissible and implementable matrix $\mathbf{\Gamma}$ in the original learning setting, then we show how to construct an admissible and implementable matrix $\mathbf{\Gamma}_c$ for the contextual learning. For this reduction we will use the notion of a separator of the policy space, defined in [35].

Definition 6.1 (Separator)**.** *A set $S \subseteq \Sigma$ of contexts is a separator for a policy space Π, if for any two policies $\pi, \pi' \in \Pi$, there exists a context $\sigma \in S$ such that $\pi(\sigma) \neq \pi'(\sigma)$.*

Definition 6.2 (Contextual S-extension of a matrix $\mathbf{\Gamma}$)**.** *For any matrix $|\mathcal{X}| \times N$ matrix $\mathbf{\Gamma}$, we define its contextual S-extension $\mathbf{\Gamma}^S$ as an $|\Pi| \times (|S| \cdot N)$ matrix, where each column j_c is associated with a pair (σ, j) for $\sigma \in S$ and $j \in [N]$ and such that the value of $\mathbf{\Gamma}^S$ at coordinates $(\pi, (\sigma, j))$ is equal to the entry of matrix $\mathbf{\Gamma}$ at coordinates $(\pi(\sigma), j)$, i.e.,*

$$\Gamma^S_{\pi,(\sigma,j)} = \Gamma_{\pi(\sigma)j}. \tag{18}$$

Lemma 6.3 (Contextual Admissibility)**.** *If $\mathbf{\Gamma}$ is a (κ, δ)-admissible matrix for a learning setting and S is a separator for the policy Π, then the contextual S-extension $\mathbf{\Gamma}_c$ is a (κ, δ)-admissible matrix for the contextual learning setting.*

Proof. First we argue that $\mathbf{\Gamma}^S$ has distinct rows. We will show that for any two rows π, π', there exists a column (σ, j) on which they differ. Since S is a separator, there must exist one $\sigma^* \in S$ on which $\pi(\sigma^*) \neq \pi'(\sigma^*)$. Now by the admissibility of the original matrix $\mathbf{\Gamma}$, we know that for any two $x, x' \in \mathcal{X}$, there exists a column j of the original matrix such that $\Gamma_{x,j} \neq \Gamma_{x',j}$. Applying the latter for $\pi(\sigma^*), \pi'(\sigma^*) \in \mathcal{X}$, we get that there exists a j^*, such that $\Gamma_{\pi(\sigma^*),j^*} \neq \Gamma_{\pi'(\sigma^*),j^*}$. By the definition of $\mathbf{\Gamma}^c$, the latter implies that $\Gamma^c_{\pi,(\sigma^*,j^*)} \neq \Gamma_{\pi',(\sigma^*,j^*)}$. Thus the two rows π, π' of matrix $\mathbf{\Gamma}$ differ at column (σ^*, j^*).

Now we argue that the quantities κ, δ (as defined in Definition 3.2), for the matrix $\mathbf{\Gamma}^S$ are the same as the quantities for matrix $\mathbf{\Gamma}$. We remind the reader that κ is the maximum number of distinct elements that appear in a column and δ is the minimum non-zero absolute difference between any two elements of a column. The reason that these quantities don't change is that they depend only on the set of values that appear in a column and not of the actual vector of values. Moreover, the set of quantities that appear in column (σ, j) of matrix $\mathbf{\Gamma}^S$, is a subset of the set of quantities that appear in column j of matrix $\mathbf{\Gamma}$. Since by taking a subset of the column we cannot increase the number of distinct elements and the minimum non-zero absolute difference of any two elements in the set, we conclude that $\mathbf{\Gamma}^S$ is also (κ, δ)-admissible. □

Lemma 6.4 (Contextual Implementability)**.** *If a pair $(\mathbf{\Gamma}, D)$ of a translation matrix and a distribution is implementable in the original learning setting, then the pair $(\mathbf{\Gamma}^S, D)$ is implementable in the contextual learning setting.*

Proof. The intuition of the proof is as follows: if we know that we can simulate every column in $\mathbf{\Gamma}$ with a sequence of polynomially many weighted inputs $\{(y, w_y)\}$, then we can simulate each column of $\mathbf{\Gamma}_c$ associated with context $\sigma \in S$ and column i of $\mathbf{\Gamma}$, by a sequence of weighted contextual inputs $\{(\sigma, y, w_y)\}$, which is essentially a contextually annotated copy of the sequence of inputs we used to simulate column i of $\mathbf{\Gamma}$. We now show this intuition more formally.

Implementability in the original setting implies that for any column $i \in \{1, \ldots, N\}$ of matrix $\mathbf{\Gamma}$ there exists a constant c_i such that for any $\alpha \in \text{supp}(D)$ there exist weights $\{w_y^{i,\alpha}\}_{y \in \mathcal{Y}}$ such that $|\{y : w_y^{i,\alpha} > 0\}| \leq \text{poly}(N, T)$ and $\alpha\left(\Gamma_{xi} - c_i\right) = \sum_{y \in \mathcal{Y}} w_y^{i,\alpha} f(x, y)$ and for all $x \in \mathcal{X}$.

In the contextual learning setting we need to argue that for any column $(\sigma, i) \in S \times [N]$ of matrix $\mathbf{\Gamma}^S$ there exists a constant $r_{(\sigma,i)}$ such that for any $\alpha \in \text{supp}(D)$ there exist weights $\{w^{(\sigma,i),\alpha}_{(\sigma',y)}\}_{(\sigma',y) \in \Sigma \times \mathcal{Y}}$, such that $|\{(\sigma', y) : w^{(\sigma,i),\alpha}_{(\sigma',y)} > 0\}| \leq poly(N, T)$ and for all $\pi \in \Pi$

$$\alpha(\Gamma^c_{\pi,(\sigma,i)} - r_{(\sigma,i)}) = \sum_{(\sigma',y) \in \Sigma \times \mathcal{Y}} w^{(\sigma,i),\alpha}_{(\sigma',y)} f_c(\pi, (\sigma', y)) \equiv \sum_{(\sigma',y) \in \Sigma \times \mathcal{Y}} w^{(\sigma,i),\alpha}_{(\sigma',y)} f(\pi(\sigma'), y).$$

The construction of such weights is straightforward: for each $y \in \mathcal{Y}$, set $w^{(\sigma,i),\alpha}_{(\sigma,y)} = w^{i,\alpha}_y$ and for $\sigma' \neq \sigma$ set $w^{(\sigma,i),\alpha}_{(\sigma',y)} = 0$. Moreover, we set $r_{(\sigma,i)} = c_i$ Now observe that:

$$\alpha(\Gamma^c_{\pi,(\sigma,i)} - r_{(\sigma,i)}) = \alpha(\Gamma_{\pi(\sigma),i} - c_i) = \sum_{y\in\mathcal{Y}} w^{i,\alpha}_y f(\pi(\sigma), y) = \sum_{(\sigma',y)\in\Sigma\times\mathcal{Y}} w^{(\sigma,i),\alpha}_{(\sigma',y)} f(\pi(\sigma'), y).$$

The latter completes the proof. □

Given the above two lemmas we can now invoke Theorem 3.5 applied to the contextual setting and get the result:

Corollary 6.5. *Let $\mathbf{\Gamma}$ be a (κ,δ)-admissible matrix with N columns and $\gamma = \max_{i,j}\Gamma_{ij}$ and let D be the uniform distribution on $[0, 1/\eta]$ for $\eta = \sqrt{\gamma\delta/(2T\kappa)}$. Then the regret of the Generalized FTPL algorithm applied to the contextual setting and with the contextual extension $\mathbf{\Gamma}^S$ of matrix $\mathbf{\Gamma}$ with respect to a separator S of size d, achieves regret:*

$$\textsc{Regret} \leq 2Nd\sqrt{2T\kappa\gamma/\delta}.$$

Moreover, assuming access to a contextual offline oracle, the latter algorithm can be implemented efficiently in time $\mathrm{poly}(N, d, T)$.

In the context of online binary classification where the features are vectors of boolean variables, [35] presents a set of examples of policy spaces that are separable. To portray the applicability of our contextual theorem, we present an example of a policy space for an auction problem that is also separable.

Example 6.6 (ORs of Boolean Features)**.** *Consider the case where each feature vector σ is a boolean vector in $\{0,1\}^r$. Moreover, consider the case of online optimal auction design over single-item second price auctions with bidder specific reserves. An example of a policy space that is separable is a policy that resembles an A/B testing over reserve prices. The auctioneer splits his reserve prices into two disjoint sets A and B (for instance, A could be a set of vectors of low reserve prices, below some threshold, and B is a set of vectors of high reserves prices). Then the set of policies is as follows: for each context $\sigma = (\sigma_1, \ldots, \sigma_r)$ take the OR of a subset of the boolean coordinates. If the OR evaluates to* 1*, then pick some vector of bidder specific reserve prices from set A and if the OR is* 0 *then pick any vector of reserve prices from set B. The whole set of policies is then generated by picking which subset of the coordinates to use for the OR function and which vector of reserves from A to output on* 1 *and which vector of reserves from B to output on* 0*. By running our contextual online learning algorithm over such a policy space, the auctioneer can learn which of the boolean coordinates are indicative of a high valuation for the players and when any one of them is* 1 *it will place a high reserve price, optimizing also over the exact vector of reserves within a high regime. Otherwise it will find an optimal vector of reserves from a low regime.*

Observe that the set of r boolean vectors that are 1 *on only one coordinate, plus the all-zero boolean vector, is a separator for this policy space. Consider any two set of different policies π, π'. If they use the same OR function, then they must differ in which vector of reserves they use from A or B. Thus they must be outputting a different action at some of the boolean vectors of the separator. If they differ in the OR function, then suppose that one includes coordinate j and the other doesn't. Then on the boolean vector which is* 1 *on only coordinate j, the one policy evaluates to* 1 *and picks a reserve vector from set A, while the other evaluates to* 0 *and picks a reserve vector from set B which has to be different, since A and B are disjoint. Thus the size of the separator for this class of policies is of order $O(r)$, while the size of the policy space is exponential in r.*

Similar analysis can be done when the policy is using the AND of boolean features rather than the OR.

6.2 Transductive Contextual Setting

We now consider the case where the learner knows a priori the set of contexts that could potentially arise, i.e., we assume that he knows that all contexts $\sigma_1, \ldots, \sigma_T$ come from a set S, which we will refer to as the *transductive set*. We do not require that he knows the multiplicity of each context or the sequence under which contexts arrive. However, we do allow the set S to be of size as large as T.

In this setting, ignoring some technical details, we could treat the transductive set S as a separator set from the last section. However, by doing so, the regret guarantee that we would derive from the analysis of the previous section grows super-linearly in T, when the size of the set S is of $\Theta(T)$. Thus in order to guarantee sub-linear regret in the transductive setting we need a tighter analysis. To achieve this we will leverage the fact that in the transductive setting S is the whole set of contexts that will arise, which is a stronger property than being a separator. This will allow us to prove a stronger stability lemma than the one that would arise from invoking Lemma 3.4 for the contextual learning setting. Moreover, we will use a perturbation distribution D that has mean zero, leading to cancellations in the cumulative error term, that make that term not grow linearly with the size of S but rather as the square root of the size of S. The combination of these two improvements leads to sub-linear regret even when $|S| = \Theta(T)$.

We will still assume that we have a (κ, δ)-admissible translation matrix $\mathbf{\Gamma}$ for the non-contextual problem and a $(\rho, 1/\delta)$-dispersed distribution D, such that $(\mathbf{\Gamma}, D)$ is implementable in the non-contextual setting. The algorithm that we will analyze is the Generalized FTPL algorithm in the contextual learning setting, with the contextual extension $\mathbf{\Gamma}^S$ of matrix $\mathbf{\Gamma}$ (where S is the transductive set) and with the same distribution D.

We first show an improved stability property of this algorithm. Even though the number of columns of matrix $\mathbf{\Gamma}^S$ is $|S| \cdot N$ we can show that the stability of the algorithm does not depend on $|S|$.

Lemma 6.7. *Consider the Generalized FTPL algorithm in the contextual learning setting with translation matrix the contextual extension $\mathbf{\Gamma}^S$ of a (κ, δ)-admissible matrix $\mathbf{\Gamma}$ with N columns and with a $(\rho, 1/\delta)$-dispersed distribution D. Then,*

$$\mathbb{E}\left[\sum_{t=1}^{T} f_c(\pi_{t+1}, (\sigma_t, y_t)) - f_c(\pi_t, (\sigma_t, y_t))\right] \leq 2TN\kappa\rho.$$

Proof. We will show that for each $t \leq T$,

$$\mathbb{E}[f_c(\pi_{t+1}, (\sigma_t, y_t)) - f_c(\pi_t, (\sigma_t, y_t))] \leq 2N\kappa\rho. \tag{19}$$

Since $f_c(\pi, (\sigma_t, y_t)) = f(\pi(\sigma_t), y_t)$ and since $f(x, y) \in [0, 1]$, it suffices to show that $\Pr[\pi_{t+1}(\sigma_t) \neq \pi_t(\sigma_t)] \leq 2N\kappa\rho$.

Observe that if $\pi_{t+1}(\sigma_t) \neq \pi_t(\sigma_t)$ then it must be that $\Gamma_{\pi_{t+1}(\sigma_t),i} \neq \Gamma_{\pi_t(\sigma_t),i}$, for some $i \in [N]$, by the admissibility of matrix $\mathbf{\Gamma}$ in the non-contextual setting. Thus we need to show that the probability that there exists an i such that $\Gamma_{\pi_{t+1}(\sigma_t),i} \neq \Gamma_{\pi_t(\sigma_t),i}$ is at most $2N\kappa\rho$. By the union bound it suffices to show that for any given i, the probability that $\Gamma_{\pi_{t+1}(\sigma_t),i} \neq \Gamma_{\pi_t(\sigma_t),i}$ is at most $2\kappa\rho$.

The proof of this fact then follows identical arguments as in the proof of the non-contextual stability Lemma 3.4 and we omit it for succinctness. □

Finally, we also show that if one uses a mean-zero Gaussian distribution D with sufficiently high variance, then the error term that is accumulated grows as the square root of the number of columns of the contextual matrix, rather than linearly. The proof of this requires a more detailed analysis of the Gaussian noise, and can be found in the appendix. Combining the two improvements we get the following theorem (see Appendix F for the proof).

Theorem 6.8. *Let $\mathbf{\Gamma}$ be a (κ, δ)-admissible matrix with N columns and $\gamma = \max_{i,j} \Gamma_{ij}$ and let $D = U[-\nu, \nu]$. Then the regret of the Generalized FTPL algorithm in the transductive contextual learning setting with translation matrix $\mathbf{\Gamma}^S$ and with the same distribution D can be bounded as*

$$\text{REGRET} \leq \frac{2TN\kappa}{\nu\delta} + 2\gamma\nu\sqrt{2N|S|\ln|\mathcal{X}|}.$$

For $\nu = \frac{\sqrt{T\kappa}N^{\frac{1}{4}}}{\sqrt{\delta\gamma}(2|S|\ln|\mathcal{X}|)^{\frac{1}{4}}}$ the above bound becomes $O\left(N^{\frac{3}{4}}(|S|\ln|\mathcal{X}|)^{\frac{1}{4}}\sqrt{\frac{T\kappa\gamma}{\delta}}\right)$. Since $|S| \leq T$, the latter is $O\left((NT)^{\frac{3}{4}}(\ln|\mathcal{X}|)^{\frac{1}{4}}\sqrt{\frac{\kappa\gamma}{\delta}}\right)$.

Observe that our application on VCG auctions with bidder-specific reserves and on envy-free item pricing continue to hold even for perturbation distributions D that have support on the negative part of the real line. Thus the results for those applications immediately extend to the transductive contextual setting by invoking the theorem above.

7 Approximate Oracles and Approximate Regret

As we saw, the Oracle-Based Generalized FTPL algorithm requires an oracle to choose the optimal learner's action x_t at time step t against a set of adversary's actions. In this section, we extend our main results regarding the Oracle-Based Generalized FTPL algorithm to work with oracles that only return *approximately optimal* actions for the learner. An offline approximation oracle is defined formally as follows:

Definition 7.1 (Offline Approximation Oracle)**.** *An* offline approximation oracle *for set of learner's actions $\mathcal{X}$ and function f, C-$\text{OPT}_{(f,\mathcal{X})}$ where $C \leq 1$, is any algorithm that receives as input a weighted set of adversary actions $S = \{(w_\ell, y_\ell)\}_{\ell \in \mathcal{L}}$, $w_k \in \mathbb{R}^+$, $y_k \in \mathcal{Y}$, and returns $x \in \mathcal{X}$, such that*

$$\sum_{(w,y)\in S} wf(x,y) \geq C \max_{x\in\mathcal{X}} \sum_{(w,y)\in S} wf(x,y).$$

As discussed earlier, access to a $(1-\epsilon)$-approximate oracle is needed for achieving no-regret results. That is, using standard online-to-batch reductions [6, 11], one can turn a polynomial time online no-regret algorithm into a polynomial time additive approximation scheme for the offline problem. So, when the best approximation for a problem is obtained through a C-optimization oracle, there is no hope for achieving no-regret results. Instead Kakade et al. [26] introduced an alternative measure of regret, called C-REGRET, for competing with offline approximation algorithm. Formally, the C-REGRET of an online maximization problem is defined as

$$C\text{-REGRET} = C \max_{a\in\mathcal{A}} \frac{1}{T}\sum_{t=1}^{T} f(x, y_t) - \mathbb{E}\left[\frac{1}{T}\sum_{t=1}^{T} f(x_t, y_t)\right].$$

We consider several types of approximation oracles, such as $(1-\epsilon)$-approximate oracle (e.g., FPTAS algorithms), approximation through relaxation of the objective, and approximation oracles that stem from Maximal-in-Range (MIR) algorithms. We show that a natural extension of the Oracle-Based Generalized FTPL algorithm,, with approximation oracle, displayed below, obtains vanishing REGRET or C-REGRET.

$(1-\epsilon)$-Approximation Oracle

We show that any FPTAS approximation algorithm is sufficient for achieving the main results of our paper. In particular, we can extend the analysis of Theorem 3.5, and the stability lemma (Lemma 3.1) to work with $(1-\epsilon)$-approximation while incurring an additional regret of ϵT. For the choice of $\epsilon = \mathcal{O}(1/\sqrt{T})$ this is truly a no-regret algorithm. This theorem is also the basis for handling working with the weaker integral oracle and its proof is deferred to section C.

Algorithm 3: Approximate Oracle-Based Generalized FTPL

Input: positive $|\mathcal{X}| \times N$ and an implementable $(\mathbf{\Gamma}, D)$ and an offline approximation oracle C-OPT
Draw $\alpha_j \sim D$ for $j = 1, \ldots, N$.
for $t = 1, \ldots, T$ **do**
  Set $S \leftarrow \{(1, y_1), \ldots, (1, y_{t-1})\} \cup \bigcup_{j \le N} S_j(\alpha_j)$.
  Set $x_t = C\text{-OPT}_{(f,\mathcal{X})}(S)$.
  Observe y_t and receive payoff $f(x_t, y_t)$.
end for

Theorem 7.2. *Let $\mathbf{\Gamma}$ be a (κ, δ)-admissible matrix with N columns and $\gamma = \max_{x,j} \Gamma_{xj}$ and let D be the uniform distribution on $[0, 1/\eta]$ for $\eta = \sqrt{\gamma\delta/(2T\kappa)}$. Furthermore, assume $(\mathbf{\Gamma}, D)$ is implementable with complexity $g(N, T)$. Then, Algorithm 3 with a $(1-\epsilon)$-OPT oracle is oracle-efficient with per-round complexity $\mathcal{O}(T + g(N, T))$ and regret of*

$$\text{REGRET} \le 2N\sqrt{2T\kappa\gamma/\delta} + \epsilon T.$$

Approximation through Relaxation

Another large class of approximation algorithms achieve their approximation guarantees by (exactly) optimizing a relaxation of the objective functions. More formally, if there is a function $F : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$, such that $Cf(x, y) \le F(x, y) \le f(x, y)$ and there is an offline approximation oracle $\text{OPT}_{(F,\mathcal{X})}$, then any online algorithm for F is also online algorithm with vanishing C-REGRET for f. This result is more formally stated below.

Theorem 7.3. *Let F be a functions such that for any $x \in \mathcal{X}$ and $y \in \mathcal{Y}$, $f(x, y) \ge F(x, y) \ge Cf(x, y)$, Let $\mathbf{\Gamma}^F$ be a (κ, δ)-admissible matrix with N columns and $\gamma = \max_{x,j} \Gamma_{xj}$ and let D be the uniform distribution on $[0, 1/\eta]$ for $\eta = \sqrt{\gamma\delta/(2T\kappa)}$. Furthermore, assume $(\mathbf{\Gamma}^F, D)$ is implementable for function F with complexity $g(N, T)$ then Algorithm 3 is oracle-efficient with per-round complexity $\mathcal{O}(T + g(N, T))$ and regret of*

$$C\text{-REGRET} \le 2N\sqrt{2T\kappa\gamma/\delta}.$$

A similar observation was made by Balcan and Blum [2] regarding approximation algorithms that use linear optimization as a relaxation and therefore can be efficiently optimized by the standard FTPL algorithm of Kalai and Vempala [27]. Our work extends this observation to any relaxation of function f that has an FPTAS and an admissible and implementable translation matrix.

Roughgarden and Wang [33] as a Relaxation. The approach of Roughgarden and Wang [33] for achieving a $1/2$-regret for single-item second price auctions with bidder-specific reserves, falls exactly in the relaxation approximation framework. They give a relaxed objective which admits a polynomial time offline oracle and which is always within a factor 2 from the original objective. Then they run an oracle based online learning algorithm for the relaxed objective. However, in their case the relaxed objective corresponds to an online linear optimization problem and can be solved with the standard FTPL algorithm of Kalai and Vempala [27]. The theorem above shows that the same approach extends even if the relaxed objective does not reduce to an online linear optimization problem but to a problem that can be tackled by our Generalized FTPL.

Approximation by Maximal-in-Range Algorithms

Another interesting class of approximation algorithms is Maximal-in-Range (MIR) algorithms. An MIR algorithm commits to a set of feasible solutions $\mathcal{X}' \subseteq \mathcal{X}$ independently of the input to the algorithm and

outputs the best solution $x \in \mathcal{X}'$. That is, an MIR C-approximation algorithm forms an approximation oracle $C\text{-}\textsc{Opt}_{(f,\mathcal{X})}(S) = \textsc{Opt}_{(f,\mathcal{X}')}(S)$ for any S. Consider an MIR approximation algorithm and $\mathbf{\Gamma}^{\mathcal{X}}$ translation matrix and distribution D such that $\mathbf{\Gamma}^{\mathcal{X}}$ is admissible for $\mathcal{X}$ and $(\mathbf{\Gamma}^{\mathcal{X}}, D)$ is implementable. Clearly, $\mathbf{\Gamma}^{\mathcal{X}}$ restricted to the set of rows in $\mathcal{X}'$ is also admissible and implementable for any $\mathcal{X}' \subseteq \mathcal{X}$. In fact $\mathbf{\Gamma}^{\mathcal{X}'}$ is (κ', δ') admissible for $\kappa' \leq \kappa$ and $\frac{1}{\delta'} \leq \frac{1}{\delta}$. Thus even better regret guarantees could be achievable if one uses the smaller admissibility quantities of matrix $\mathbf{\Gamma}^{\mathcal{X}'}$. Therefore, an MIR C-approximation algorithm leads to an efficient online algorithm with vanishing C-$\textsc{Regret}$. More formally we have:

Theorem 7.4. *Let $\mathbf{\Gamma}^{\mathcal{X}}$ be a (κ, δ)-admissible matrix with N columns and $\gamma = \max_{x,j} \Gamma_{xj}$ and let D be the uniform distribution on $[0, 1/\eta]$ for $\eta = \sqrt{\gamma\delta/(2T\kappa)}$. Furthermore, assume $(\mathbf{\Gamma}^{\mathcal{X}}, D)$ is implementable with complexity $g(N, T)$. Then Algorithm 3 used with any MIR approximation algorithm C-$\textsc{Opt}$ is an efficient algorithm with per-round complexity $\mathcal{O}(T + g(N, T))$ and regret of*

$$C\text{-}\textsc{Regret} \leq 2N\sqrt{2T\kappa\gamma/\delta}.$$

Furthermore, setting the oracle specific, $\eta = \sqrt{\gamma\delta'/(2T\kappa')}$, gives a stronger regret bound of $2N\sqrt{2T\kappa'\gamma/\delta'}$.

8 Additional Applications and Connections

In this section, we discuss an application of our oracle efficient learning approach to the problem of online welfare maximization in multi-unit auctions, to no-regret learning in simultaneous second price auctions and discuss the connections between our notions of admissibility and implementability with other statistical measures of complexity from learning theory.

8.1 Polynomial Algorithm for Online Welfare Maximization in Multi-Unit Auctions

In this section, we consider single-item multi-unit auctions that has $1/2$-approximation Maximal-in-Range (MIR) algorithm. We show how the results of Section 7 can be applied to this problem to achieve a *truly polynomial time* online algorithm with vanishing $\frac{1}{2}$-$\textsc{Regret}$.

We consider an online learning variant of an n-bidder *multi-unit* environment, better modeled as a set of k identical items. Each bidder i has a monotonic valuation function $v_i : [k] \to [0, 1]$. In other words, bidder i has *marginal valuation* $\mu_i(\ell)$ for receiving the ℓ^{th} item and the total utility of bidder i for receiving the q_ith items is $v_i(q_i) = \sum_{s=1}^{q_i} \mu_i(s)$. The goal is to find an allocation $\mathbf{q}$ that maximized the total welfare in time polynomial in n and $\log(k)$. In the online learning setting, every day t a fresh set of bidders arrives with some new valuation v_i^t and the learner commits to an allocation of the units to the players, prior to seeing the valuations. The goal of the learner is to pick an allocation each day that competes with the best in hindsight.

It is not hard to see that the offline welfare maximization problem in this setting corresponds to the Knapsack problem, where each player has a valuation equal to the average value in hindsight, i.e. $\frac{1}{t}\sum_{\tau=1}^{t} v_i^\tau(\cdot)$. So, dynamic programming can be used to compute a welfare-maximizing allocation in time polynomial in n and k. Dobzinski and Nisan [13] introduced a $1/2$-approximation MIR algorithm for this problem. At a high level, the algorithm proceeds by dividing the set of items to k/n^2 bundles of the same size.[5] Then the MIR algorithm chooses the best allocation from the set of allocations (range of the algorithm) where all the items in one bundle are allocated to the same bidders. This algorithm is effectively solving a knapsack problem over n^2 items and can be implemented in time polynomial in n and $\log(k)$.

We show how to construct a matrix $\mathbf{\Gamma}^{\text{MU}}$ that is admissible and implementable for unrestricted n-bidder k-unit auctions. We then use the theorem 7.4 from Section 7 to obtain an online algorithm with vanishing $\frac{1}{2}$-$\textsc{Regret}$ that runs in time $\text{poly}(n, T, \log k)$.

Let $\mathbf{\Gamma}^{\text{MU}}$ be a matrix with n columns, such that for any allocation $\mathbf{q}$ and any column j, $\Gamma^{\text{MU}}_{\mathbf{q},j} = q_j/k$. Clearly, for any $\mathbf{q} \neq \mathbf{q}'$ we have $\mathbf{\Gamma}^{\text{MU}}_{\mathbf{q}} \neq \mathbf{\Gamma}^{\text{MU}}_{\mathbf{q}}$. So, $\mathbf{\Gamma}^{\text{MU}}$ is $(k+1, \frac{1}{k})$-admissible. Moreover, observe that

[5] If k is not a multiple of n^2, one can add extra items with marginal value of 0 to all bidders.

the matrix $\mathbf{\Gamma}^{\mathrm{MU}'}$ restricted to the range of the $1/2$ approximation algorithm of Dobzinski and Nisan [13] has much better admissibility constants. In particular, the number of different entries within each column is at most $\kappa' = n^2 + 1$, since each player receives allocations that are multiples of k/n^2 and there are at most $n^2 + 1$ such multiples. Moreover, the minimum non-zero difference of the column entries, between any two such bundled allocations, is at least $\delta' \geq \frac{k}{n^2}\frac{1}{k} = \frac{1}{n^2}$. Therefore matrix $\mathbf{\Gamma}^{\mathrm{MU}'}$ is $\left(n^2+1, \frac{1}{n^2}\right)$-admissible.

Let D be any distribution over $[0, \alpha_{max}]$. It is not hard to see that $(\mathbf{\Gamma}^{\mathrm{MU}}, D)$ is also implementable. For any column j, consider the bid profile $\mathbf{v}^j$ where bidder j's marginal valuation for any item is $1/k$ and all other bidders have 0 valuation for any number of items. That is, $\mu_j(s) = 1/k$ and $\mu_i(s) = 0$ for all s and $i \neq j$. The welfare of any allocation on this bid profile is the utility of bidder j, which is the $q_j/k = \mathbf{\Gamma}^{\mathrm{MU}}_{\mathbf{q}}$. Therefore, $(\mathbf{\Gamma}^{\mathrm{MU}}, D)$ is implementable using the map $S_j(\alpha_j) = \{(\alpha_j, \mathbf{v}^j)\}$. Therefore, also $(\mathbf{\Gamma}^{\mathrm{MU}'}, D)$ is implementable.

The next theorem is a direct consequence of the above discussion.

Theorem 8.1. *Consider the problem of welfare maximization in k-unit auctions. Let D be the uniform distribution on $[0, 1/\eta]$ for $\eta = \sqrt{\delta'/(2T\kappa')} = \sqrt{\frac{1}{2Tn^2(n^2+1)}}$. Then, for any sequence of valuation functions $\mathbf{v}^1, \ldots, \mathbf{v}^T$, the Generalized FTPL algorithm with matrix Γ^{MU}, distribution D and oracle the $1/2$-aproximate MIR algorithm of [13], runs in per-round time $\mathrm{poly}(n, T)$, plays the sequence of allocations $\mathbf{q}^1, \ldots, \mathbf{q}^T$, and has $1/2$-*REGRET

$$\frac{1}{2}\text{-REGRET} = \frac{1}{2}\left(\frac{1}{T}\max_{\mathbf{q}\in\mathbb{Z}^n}\sum_{t=1}^{T}\sum_{i=1}^{n} v_i^t(q_i)\right) - \mathbb{E}\left[\frac{1}{T}\sum_{t=1}^{T}\sum_{i=1}^{n} v_i^t(q_i^t)\right] \leq n\cdot 2n\sqrt{2T(n^2+1)n^2} \leq O(n^4\sqrt{T}).$$

The extra factor of n in the latter bound as compared to that implied by Theorem 7.4, is due to the fact that the maximum welfare in this problem is upper bounded by n and not 1. So we need to scale-down the valuations before applying the theorem by dividing by n and then scale up again to get the regret for the un-normalized problem.

8.2 Oracle Efficient No-Regret Learning in Simultaneous Second Price Auctions

In this section, we answer an open problem raised by Daskalakis and Syrgkanis [11] regarding the existence of an oracle-based no-regret algorithm for optimal bidding in Simultaneous Second-Price Auctions. We show that our Oracle-Based Generalized FTPL algorithm used with an appropriate implementable and admissible translation matrix can be used to obtain such an algorithm.

A *Simultaneous Second-Price Auction (SiSPA)* [3, 9, 14] is a mechanism for allocating k items to n bidders. Each bidder $i \leq n$ submits k simultaneous bids denoted by a vector of bids $\mathbf{b}_i$. The mechanism allocates each item using a second-price auction based on the bids solely submitted for this item, while breaking ties in favor of bidders with larger indices. For each item j, the winner is charged θ_j, the second highest bid for that item. Each bidder i has a fixed combinatorial valuation function $v_i : \{0,1\}^k \to [0,1]$ over bundles of items. Then, the total utility of bidder i who is allocated the bundle $\mathbf{q}_i \in \{0,1\}^k$ is $v_i(\mathbf{q}_i) - \boldsymbol{\theta}\cdot\mathbf{q}_i$, where $\boldsymbol{\theta}$ is the vector of second largest bids across all items.

We consider the problem of optimal bidding in a SiSPA from the perspective of the first bidder. Hereafter, we drop the indices of the players from the notation. From this perspective, the utility of the bidder only depends on its bid $\mathbf{b}$ and the the threshold vector $\boldsymbol{\theta}$ of the second largest bids. The online optimal bidding problem is defined as follows.

Definition 8.2 (Online Bidding in SiSPAs [35])**.** *At each time step t, the player picks a bid vector $\mathbf{b}^t$ and an adversary picks a threshold vector $\boldsymbol{\theta}^t$. The player wins the bundle of items $\mathbf{q}(\mathbf{b}^t, \boldsymbol{\theta}^t)$, with $q_j(\mathbf{b}^t, \boldsymbol{\theta}^t) = \mathbf{1}_{(b_j^t > \theta_j^t)}$ and gets the utility*

$$u(\mathbf{b}^t, \boldsymbol{\theta}^t) = v\left(\mathbf{q}(\mathbf{b}^t, \boldsymbol{\theta}^t)\right) - \boldsymbol{\theta}^t\cdot\mathbf{q}.$$

We consider this problem under the *no-overbidding condition* that requires that for any bundle $\mathbf{q}$, the sum of bids over items in $\mathbf{q}$ does not exceed the bidder's valuation for $\mathbf{q}$, i.e., $\mathbf{b}^t \cdot \mathbf{q} \leq v(\mathbf{q})$, for all $\mathbf{q} \in \{0,1\}^k$. Similar no-overbidding assumptions are used in the previous work to prove that no-regret learning in second-price auctions has good welfare guarantees [9, 14].

We consider the online bidding problem where for any $\mathbf{q}$, the valuation $v(\mathbf{q})$ is a multiple of $1/m$ and for any item $j \leq k$, b_j is a multiple of $1/m$. We represent by $\mathcal{B}_m$ the class of all such bid vectors that satisfy the no-overbidding condition for $v(\cdot)$. Note that the assumption on the valuation function is not restrictive. That is, for any valuation function $v(\cdot)$, one can round down its values to the closest multiple of $1/m$, while losing at most $1/m$ utility. Moreover, a similar discretization for the bid vectors was used by Daskalakis and Syrgkanis [35] for studying offline and online optimal bidding in SiSPAs.

Next, we show how to construct an implementable and admissible translation matrix for $\mathcal{B}_m$. Let $\mathbf{\Gamma}^{\text{IB}}$ be a matrix with k columns and $|\mathcal{B}_m|$ rows that are equal to bid vectors, i.e., $\mathbf{\Gamma}^{\text{IB}}_{\mathbf{b}} = \mathbf{b}$. This immediately yields admissibility with $\kappa = m+1$ and $\delta = 1/m$, because of the discretization of bids. Next, we show that for any D with a bounded non-negative support, $(\mathbf{\Gamma}^{\text{IB}}, D)$ is also implementable. At a high level, to implement the jth perturbation, we consider threshold vectors in which $\theta_j = \ell/m$ for $\ell = 0, 1, \ldots, m-1$, and all other thresholds are set to the highest level. In such threshold vectors, the bidder can win at most one item, item j. Furthermore, whenever the bidder wins item j, its utility is independent of the bid vector, as it only depends on the valuation $v(\mathbf{e}_j)$ and the entry θ_j of the threshold vector. So, there are weights such that the total contribution of weighted threshold vectors towards the utility of any *winning bid* is the same across all ℓ. Since bid b_j wins item j for all $\ell/m < b_j$, then the total utility of playing b_j against this weighted set of threshold vectors is proportional to $b_j = \Gamma^{\text{IB}}_{\mathbf{b}j}$. This implies that $(\mathbf{\Gamma}^{\text{IB}}, D)$ is implementable. The next lemma formalizes this discussion.

Lemma 8.3. *Let D be a distribution supported on a subset of $[0, \alpha_{\max}]$. Then $(\mathbf{\Gamma}^{\text{IB}}, D)$ is implementable with complexity $\text{poly}(\alpha_{\max}, m, k)$, and $\mathbf{\Gamma}^{\text{IB}}$ is $(m+1,\ 1/m)$-admissible.*

Proof. We argued the admissibility above, so it suffices to show implementability. The dataset $S_j(\alpha_j)$ will contain threshold vectors where all but the jth thresholds are set to 1. Specifically, for $\ell = 0, 1, \ldots, m-1$, let $\boldsymbol{\theta}_\ell = (\ell/m)\mathbf{e}_j + \sum_{j' \neq j} \mathbf{e}_{j'}$. Note that the utility of playing bid $\mathbf{b}$ against $\boldsymbol{\theta}_\ell$ is $u(\mathbf{b}, \boldsymbol{\theta}_\ell) = \big(v(\mathbf{e}_j) - \ell/m\big)\mathbf{1}_{(b_j > \ell/m)}$. We set the weight corresponding to $\boldsymbol{\theta}_\ell$ to

$$w_\ell = \begin{cases} \frac{\alpha_j}{m} \cdot \frac{1}{v(\mathbf{e}_j) - \ell/m} & \text{if } \ell/m < v(\mathbf{e}_j) \\ 0 & \text{otherwise.} \end{cases}$$

Since $b_j \leq v(\mathbf{e}_j)$ for any $\mathbf{b}$, we have

$$\begin{aligned} \sum_{\ell=0}^{m-1} w_\ell\, u(\mathbf{b}, \boldsymbol{\theta}_\ell) &= \sum_{\ell=0}^{m-1} \frac{\alpha_j}{m} \cdot \frac{1}{v(\mathbf{e}_j) - \ell/m} \cdot \Big(v(\mathbf{e}_j) - \ell/m\Big)\mathbf{1}_{(b_j > \ell/m)} \\ &= \sum_{\ell=0}^{m-1} \frac{\alpha_j}{m} \mathbf{1}_{(b_j > \ell/m)} = \sum_{\ell=0}^{mb_j - 1} \frac{\alpha_j}{m} \\ &= \alpha_j b_j = \alpha_j \Gamma^{\text{IB}}_{\mathbf{b}j}. \end{aligned}$$

Thus, indeed $\{(w_\ell, \mathbf{v}_\ell)\}_\ell$ implements α_j. This set can be clearly generated in time $\text{poly}(k, m)$, its cardinality is m and its sum of weights is at most $\alpha_{\max} m$, since each weight w_ℓ is at most $\alpha_{\max}$. Therefore, $(\mathbf{\Gamma}^{\text{IB}}, D)$ is implementable with complexity $\text{poly}(k, m, \alpha_{\max})$. □

The next theorem is a direct consequence of Lemma 8.3.

Theorem 8.4. *Consider the problem of* Online Bidding in SiSPAs. *Let D be the uniform distribution as described in Theorem 3.5. Then, the Oracle-Based Generalized FTPL algorithm with $\mathbf{\Gamma}^{\text{IB}}$ and D is oracle-efficient with complexity* $\text{poly}(k, m, T)$ *and has regret*

$$\mathbb{E}\left[\max_{\mathbf{b}\in\mathcal{B}_m}\sum_{t=1}^{T}u(\mathbf{b},\boldsymbol{\theta}^t)-\sum_{t=1}^{T}u(\mathbf{b}^t,\boldsymbol{\theta}^t)\right]\leq O(km\sqrt{T}).$$

8.3 Universal Identification Sequences

There is an interesting connection between our definitions of admissibility and implementability and a statistical measure of complexity from learning theory, called the *Universal Identification Sequences*.

Definition 8.5 (Universal Identification Sequences [16])**.** *Consider a domain $\mathcal{Z}$ and a class of functions $\mathcal{F}$ such that for all $f\in\mathcal{F}$, $f:\mathcal{Z}\to\{0,1\}$. A sequence of unlabeled instances $S=\{z\}$ is said to* distinguish *function $f\in\mathcal{F}$ if f is the only function that is consistent with the labeling on S produced by f. A sequence of unlabeled instances $S=\{z\}$ is called a* universal identification sequence *if it* distinguishes *every $f\in\mathcal{F}$.*

Any universal identification sequence of $\mathcal{F}$ can be used to construct a translation matrix that is admissible and implementable. Consider a matrix $\mathbf{\Gamma}^{\mathcal{F}}$, whose rows are indexed by $\mathcal{F}$ and columns are indexed by S, such that $\mathbf{\Gamma}^{\mathcal{F}}_{fz}=f(z)$ for any $f\in\mathcal{F}$ and $z\in S$. By the definition of universal identification sequence for any two functions, $f,f'\in\mathcal{F}$ there is $z\in S$, such that $f(z)\neq f'(z)$, i.e., $\mathbf{\Gamma}^{\mathcal{F}}_{f}\neq\mathbf{\Gamma}^{\mathcal{F}}_{f'}$. As a results $\mathbf{\Gamma}^{\mathcal{F}}$ is $(2,1)$-admissible. Moreover, the columns of $\mathbf{\Gamma}^{\mathcal{F}}$ correspond to the value of functions applied to $z\in S$. Therefore, $(\mathbf{\Gamma}^{\mathcal{F}},D)$ is implementable for any D over $\mathbb{R}^+$. That is, the *length of a universal identification sequence* is an upper bound on the number of columns needed to create a translation matrix that is admissible and implementable for a class of binary functions. Examples of function classes with polynomial-length universal identification sequences include logarithmic-depth read-once majority formulas and logarithmic-depth read-once positive NAND formulas [16]. The next corollary is a direct consequence of Theorems 3.5 and Theorems 3.11.

Corollary 8.6. *Consider a domain $\mathcal{Z}$ and a class of binary functions $\mathcal{F}$ with a universal identification sequence of length d. Then, there is an efficient oracle-based online algorithm with regret*

$$\mathbb{E}\left[\max_{f\in\mathcal{F}}\sum_{t=1}^{T}f(z_t)-\sum_{t=1}^{T}f_t(z_t)\right]\leq O(d\sqrt{T}).$$

Our Oracle-Based Generalized FTPL algorithm goes beyond the use of binary functions and universal identification sequences. In particular, we applied our results to obtain no-regret algorithms for several commonly studied classes of auctions (the revenue of which are real-valued functions). Furthermore, we introduced implementable translation matrices where each column corresponds to a complex weighted set of adversary's actions, rather than, columns that correspond to individual adversary's actions.

Acknowledgments

We would like to thank Avrim Blum for helpful discussions and pointers regarding the universal identification sequences.

References

[1] Awerbuch, B. and Kleinberg, R. (2008). Online linear optimization and adaptive routing. *J. Comput. Syst. Sci.*, 74(1):97–114.

[2] Balcan, M.-F. and Blum, A. (2006). Approximation algorithms and online mechanisms for item pricing. In *Proceedings of the 7th ACM Conference on Economics and Computation (EC)*, pages 29–35. ACM.

[3] Bhawalkar, K. and Roughgarden, T. (2011). Welfare guarantees for combinatorial auctions with item bidding. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 700–709. SIAM.

[4] Blackwell, D. (1956). An analog of the minimax theorem for vector payoffs. *Pacific J. Math.*, 6(1):1–8.

[5] Blum, A. and Hartline, J. D. (2005). Near-optimal online auctions. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1156–1163. Society for Industrial and Applied Mathematics.

[6] Cesa-Bianchi, N., Conconi, A., and Gentile, C. (2006). On the generalization ability of on-line learning algorithms. *IEEE Trans. Inf. Theor.*, 50(9):2050–2057.

[7] Cesa-Bianchi, N., Gentile, C., and Mansour, Y. (2013). Regret minimization for reserve prices in second-price auctions. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1190–1204. Society for Industrial and Applied Mathematics.

[8] Cheeger, J. (1969). *A lower bound for the smallest eigenvalue of the Laplacian*, pages 195–199.

[9] Christodoulou, G., Kovács, A., and Schapira, M. (2008). Bayesian combinatorial auctions. In *Proceedings of theInternational Colloquium on Automata, Languages and Programming (ICALP)*, pages 820–832. Springer.

[10] Cole, R. and Roughgarden, T. (2014). The sample complexity of revenue maximization. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing (STOC)*, pages 243–252. ACM.

[11] Daskalakis, C. and Syrgkanis, V. (2016). Learning in auctions: Regret is hard, envy is easy. In *Proceedings of the 57th Symposium on Foundations of Computer Science (FOCS)*.

[12] Devanur, N. R., Huang, Z., and Psomas, C.-A. (2016). The sample complexity of auctions with side information. In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing (STOC)*, pages 426–439. ACM.

[13] Dobzinski, S. and Nisan, N. (2007). Mechanisms for multi-unit auctions. In *Proceedings of the 8th ACM Conference on Economics and Computation (EC)*, pages 346–351. ACM.

[14] Feldman, M., Fu, H., Gravin, N., and Lucier, B. (2013). Simultaneous auctions are (almost) efficient. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing (STOC)*, pages 201–210. ACM.

[15] Freund, Y. and Schapire, R. E. (1995). A desicion-theoretic generalization of on-line learning and an application to boosting. In *European conference on computational learning theory*, pages 23–37. Springer.

[16] Goldman, S. A., Kearns, M. J., and Schapire, R. E. (1993). Exact identification of read-once formulas using fixed points of amplification functions. *SIAM Journal on Computing*, 22(4):705–726.

[17] Guruswami, V., Hartline, J. D., Karlin, A. R., Kempe, D., Kenyon, C., and McSherry, F. (2005). On profit-maximizing envy-free pricing. In *Proceedings of the16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1164–1173. Society for Industrial and Applied Mathematics.

[18] Hannan, J. (1957). Approximation to bayes risk in repeated play. *Contributions to the Theory of Games*, 3:97–139.

[19] Hartline, J. D. (2013). Mechanism design and approximation. *Book draft. October*, 122.

[20] Hartline, J. D. and Koltun, V. (2005). Near-optimal pricing in near-linear time. In *Workshop on Algorithms and Data Structures*, pages 422–431. Springer.

[21] Hartline, J. D. and Roughgarden, T. (2009). Simple versus optimal mechanisms. In *Proceedings of the 10th ACM Conference on Economics and Computation (EC)*, pages 225–234. ACM.

[22] Hazan, E. and Kale, S. (2012). Online submodular minimization. *Journal of Machine Learning Research*, 13(Oct):2903–2922.

[23] Hazan, E. and Koren, T. (2016). The computational power of optimization in online learning. In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing (STOC)*. ACM.

[24] Hutter, M. and Poland, J. (2005). Adaptive online prediction by following the perturbed leader. *Journal of Machine Learning Research*, 6:639–660.

[25] Kakade, S. and Kalai, A. T. (2005). From batch to transductive online learning. In *Proceedings of the 18th Annual Conference on Neural Information Processing Systems (NIPS)*, pages 611–618.

[26] Kakade, S. M., Kalai, A. T., and Ligett, K. (2009). Playing games with approximation algorithms. *SIAM Journal on Computing*, 39(3):1088–1106.

[27] Kalai, A. and Vempala, S. (2005). Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 71(3):291–307.

[28] Morgenstern, J. H. and Roughgarden, T. (2015). On the pseudo-dimension of nearly optimal auctions. In *Proceedings of the 29th Annual Conference on Neural Information Processing Systems (NIPS)*, pages 136–144.

[29] Myerson, R. B. (1981). Optimal auction design. *Mathematics of operations research*, 6(1):58–73.

[30] Nisan, N. and Ronen, A. (2007). Computationally feasible vcg mechanisms. *J. Artif. Int. Res.*, 29(1):19–47.

[31] Paulin, D. (2012). Concentration inequalities for Markov chains by Marton couplings and spectral methods. *ArXiv e-prints*.

[32] Roughgarden, T. and Schrijvers, O. (2016). Ironing in the dark. In *Proceedings of the17th ACM Conference on Economics and Computation (EC)*, pages 1–18. ACM.

[33] Roughgarden, T. and Wang, J. R. (2016). Minimizing regret with multiple reserves. In *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*, pages 601–616. ACM.

[34] Shalev-Shwartz, S. and Ben-David, S. (2014). *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, New York, NY, USA.

[35] Syrgkanis, V., Krishnamurthy, A., and Schapire, R. E. (2016). Efficient algorithms for adversarial contextual learning. In *Proceedings of the 33nd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, pages 2159–2168.

A Proof of Lemma 3.1

We begin by stating a variant of the standard Be-the-Leader Lemma [27] that has been modified to include the perturbation specified by the matrix Γ.

Lemma A.1 (Be-the-Leader Lemma)**.** *Under Generalized FTPL, for any $x \in \mathcal{X}$,*

$$\sum_{t=1}^{T} f(x_{t+1}, y_t) + \boldsymbol{\alpha} \cdot \boldsymbol{\Gamma}_{x_1} \geq \sum_{t=1}^{T} f(x, y_t) + \boldsymbol{\alpha} \cdot \boldsymbol{\Gamma}_x.$$

Proof. For $T = 1$ the inequality holds trivially. Assume that the claim holds for some T. Then, for all x

$$\begin{aligned}
\sum_{t=1}^{T+1} f(x_{t+1}, y_t) + \boldsymbol{\alpha} \cdot \boldsymbol{\Gamma}_{x_1} &= \sum_{t=1}^{T} f(x_{t+1}, y_t) + f(x_{T+2}, y_{T+1}) + \boldsymbol{\alpha} \cdot \boldsymbol{\Gamma}_{x_1} \\
&\geq \sum_{t=1}^{T} f(x_{T+2}, y_t) + \boldsymbol{\alpha} \cdot \boldsymbol{\Gamma}_{x_{T+2}} + f(x_{T+2}, y_{T+1}) \quad \text{(by induction hypothesis)} \\
&= \sum_{t=1}^{T+1} f(x_{T+2}, y_t) + \boldsymbol{\alpha} \cdot \boldsymbol{\Gamma}_{x_{T+2}} \\
&\geq \sum_{t=1}^{T+1} f(x, y_t) + \boldsymbol{\alpha} \cdot \boldsymbol{\Gamma}_x. \quad \text{(by optimality of } x_{T+2})
\end{aligned}$$

□

Let $x^* = \arg\max_{x \in \mathcal{X}} \sum_{t=1}^{T} f(x, y_t)$. Then by the lemma above,

$$\begin{aligned}
\textsc{Regret} &= \mathbb{E}\left[\sum_{t=1}^{T} f(x^*, y_t) - \sum_{t=1}^{T} f(x_t, y_t)\right] \\
&= \mathbb{E}\left[\sum_{t=1}^{T} f(x^*, y_t) - \sum_{t=1}^{T} f(x_{t+1}, y_t)\right] + \mathbb{E}\left[\sum_{t=1}^{T} f(x_{t+1}, y_t) - \sum_{t=1}^{T} f(x_t, y_t)\right] \\
&\leq \mathbb{E}\left[\boldsymbol{\alpha} \cdot (\boldsymbol{\Gamma}_{x_1} - \boldsymbol{\Gamma}_{x^*})\right] + \mathbb{E}\left[\sum_{t=1}^{T} f(x_{t+1}, y_t) - \sum_{t=1}^{T} f(x_t, y_t)\right].
\end{aligned}$$

B Proof of Theorem 3.11

To show that the Oracle-Based FTPL procedure (Algorithm 2) implements Generalized FTPL (Algorithm 1), it suffices to show that at each round t,

$$\arg\max_{x \in \mathcal{X}} \left[\sum_{\tau=1}^{t-1} f(x, y_\tau) + \boldsymbol{\alpha} \cdot \boldsymbol{\Gamma}_x\right] = \textsc{Opt}(S) = \arg\max_{x \in \mathcal{X}} \sum_{(w,y) \in S} w f(x, y). \tag{20}$$

Consider any $x, x' \in \mathcal{X}$. Then, from the definition of S and by implementability,

$$\sum_{(w,y) \in S} w f(x, y) - \sum_{(w,y) \in S} w f(x', y) = \sum_{\tau=1}^{t-1} \left[f(x, y_\tau) - f(x', y_\tau)\right] + \sum_{j \leq N} \sum_{(w,y) \in S_j(\alpha_j)} w \left[f(x, y) - f(x', y)\right]$$

Algorithm 4: Oracle-Based Generalized ϵ-FTPL with Integral Oracle

Input: positive $|\mathcal{X}| \times N$, an implementable $(\mathbf{\Gamma}, D)$, an oracle IntOpt, precision parameter ϵ
Draw $\alpha_j \sim D$ for $j = 1, \ldots, N$.
for $t = 1, \ldots, T$ **do**
  Set $S = \{(1, y_1), \ldots, (1, y_{t-1})\} \cup \bigcup_{j \le N} S_j(\alpha_j)$.
  Set $S' = \left\{ \left(\lfloor w|S| \,/\, \epsilon \rfloor, \, y \right) : \, (w, y) \in S \right\}$.
  Set $x_t = \textsc{IntOpt}(S')$.
  Observe y_t and receive payoff $f(x_t, y_t)$.
end for

$$= \sum_{\tau=1}^{t-1} \left[f(x, y_\tau) - f(x', y_\tau) \right] + \sum_{j \le N} \alpha_j (\Gamma_{xj} - \Gamma_{x'j})$$
$$= \left(\sum_{\tau=1}^{t-1} f(x, y_\tau) + \boldsymbol{\alpha} \cdot \mathbf{\Gamma}_x \right) - \left(\sum_{\tau=1}^{t-1} f(x', y_\tau) + \boldsymbol{\alpha} \cdot \mathbf{\Gamma}_{x'} \right),$$

which immediately yields Equation (20).

Also, by implementability, the running time to construct the set S is at most $T + g(N, T)$, and its mass is at most $T + g(N, T)$. Since there is only one oracle call per round, we get the per-round complexity of $T + g(N, T)$.

C Integral Oracles

We now study how online learning can be reduced to a more restrictive class of offline optimization oracles, which only allow the weights that are positive integers:

Definition C.1 (Integral Offline Oracle)**.** *An* integral offline oracle IntOpt *receives as input a set of adversary actions with non-negative integer weights $S = \{(w_\ell, y_\ell)\}_{\ell \in \mathcal{L}}$, $w_\ell \in \mathbb{N}$, $y_\ell \in \mathcal{Y}$, and returns*

$$\textsc{IntOpt}(S) \in \arg\max_{x \in \mathcal{X}} \sum_{(w,y) \in S} w f(x, y).$$

Integral oracles can be implemented by algorithms that optimize over unweighted data sets, by replicating each example according to its weight. A natural input size for such algorithms is the size of the data set (including the replicates). This exactly coincides with the mass of S defined in the previous section.

We will show that online learning can be efficiently reduced to offline optimization with IntOpt oracles under the same conditions as studied in the previous section. Thus, while the rest of paper assumes real-weighted oracles, the results are also valid when we only have access to an integral oracle.

At the crux of our approach is a construction of an *approximate* real-weighted oracle from an integral oracle. It turns out that such an approximate oracle, with a suitably chosen precision parameter ϵ, is all that is required for no-regret learning with the Generalized FTPL algorithm.

Lemma C.2 (Integral Offline Oracle to Approximate Offline Oracle)**.** *Given an integral offline oracle* IntOpt *and any precision parameter $\epsilon > 0$, one can construct an approximate offline oracle* Opt$'$ *which receives as input a weighted set of adversary actions $S = \{(w_\ell, y_\ell)\}_{\ell \in \mathcal{L}}$, $w_\ell \in \mathbb{R}^+$, and returns $\textsc{Opt}'(S) \in \mathcal{X}$ such that $\forall x \in \mathcal{X}$,*

$$\sum_{(w,y) \in S} w f(\textsc{Opt}'(S), y) \ge \sum_{(w,y) \in S} w f(x, y) - \epsilon.$$

This approximate oracle is implemented by a single oracle call to IntOpt(S') *on the dataset*

$$S' = \left\{ \left(\lfloor w|S|/\epsilon \rfloor, y \right) : (w, y) \in S \right\}.$$

This lemma is proved in Appendix C.1. The construction of the set S' gives rise to a variant of the Generalized FTPL algorithm, presented in Algorithm 4. This algorithm can be analyzed similarly as the Generalized FTPL to prove the regret bound essentially identical to the bound of Theorem 3.5. The algorithm is also oracle-efficient under implementability. Both results are summarized in the following theorem:

Theorem C.3. *Let* $\mathbf{\Gamma}$ *be a* (κ, δ)*-admissible matrix with* N *columns and* $\gamma = \max_{x,j} \Gamma_{xj}$ *and let* D *be the uniform distribution on* $[0, 1/\eta]$ *for* $\eta = \sqrt{\gamma\delta/(2T\kappa)}$*. If* $\epsilon = \frac{1}{\sqrt{T}}$*, then the regret of Algorithm 4 can be bounded as*

$$\textsc{Regret} \le O(N\sqrt{2T\kappa\gamma/\delta}).$$

Furthermore, if $(\mathbf{\Gamma}, D)$ *is implementable with complexity* $g(N, T)$ *then Algorithm 4 is oracle-efficient with respect to* IntOpt*, with per-round complexity* $\mathrm{poly}(T, g(N, T), 1/\epsilon)$*.*

For the proof of regret bound, see Appendix C.2. The oracle efficiency result follows from Theorem 3.11, using the fact that within the reduction from real-weighted to integral oracle, the mass of the dataset S' is bounded as $\|S'\| \le \|S\|^2/\epsilon$.

C.1 Proof of Lemma C.2

Proof of Lemma C.2. Let $\delta = \epsilon/|S|$. Thus, elements $(w, y) \in S$ are transformed into $(\lfloor w/\delta \rfloor, y) \in S'$. To show the approximate optimality, we use the fact that $0 \le a - \lfloor a \rfloor \le 1$ for any $a \ge 0$, so we have $0 \le w - \delta\lfloor w/\delta \rfloor \le \delta$ and thus for any $x \in \mathcal{X}$,

$$\begin{aligned}
\sum_{(w,y)\in S} w f(\textsc{Opt}'(S), y) &\ge \delta \sum_{(w,y)\in S} \lfloor w/\delta \rfloor f(\textsc{Opt}'(S), y) \\
&= \delta \sum_{(w',y)\in S'} w' f(\textsc{IntOpt}(S'), y) && \text{(by construction)} \\
&\ge \delta \sum_{(w',y)\in S'} w' f(x, y) && \text{(by the property of IntOpt)} \\
&\ge \sum_{(w,y)\in S} w f(x, y) - \sum_{(w,y)\in S} \delta \\
&= \sum_{(w,y)\in S} w f(x, y) - \epsilon. && \square
\end{aligned}$$

C.2 Approximate Oracles

Approximate oracles, with a precision error ϵ, give rise to the Generalized ϵ-FTPL algorithm shown in Algorithm 5. The analysis of Section 3.1 can be generalized to obtain the following result:

Theorem C.4. *Let* Γ *be a* (κ, δ)*-admissible matrix with* N *columns and* $\gamma = \max_{x,j} \Gamma_{xj}$ *and let* D *be the uniform distribution on* $[0, 1/\eta]$ *for* $\eta = \sqrt{\gamma\delta/(2T\kappa)}$*. If* $\epsilon = \frac{1}{\sqrt{T}}$*, then the regret of the Generalized* ϵ*-FTPL algorithm can be bounded as*

$$\textsc{Regret} \le O(N\sqrt{2T\kappa\gamma/\delta}).$$

The proof of this theorem is a direct application of the generalized version of Lemma 3.1 and 3.4. Details are given in the following subsections.

Algorithm 5: Generalized ϵ-FTPL

Input: non-negative matrix $\mathbf{\Gamma} \in \mathbb{R}_+^{|\mathcal{X}| \times N}$, distribution D
Draw $\alpha_j \sim D$ for $j = 1, \ldots, N$
for $t = 1, \ldots, T$ **do**
  Choose x_t such that for all $x \in \mathcal{X}$,

$$\sum_{\tau=1}^{t-1} f(x_t, y_\tau) + \boldsymbol{\alpha} \cdot \mathbf{\Gamma}_{x_t} \geq \sum_{\tau=1}^{t-1} f(x, y_\tau) + \boldsymbol{\alpha} \cdot \mathbf{\Gamma}_x - \epsilon$$

  Observe y_t and receive payoff $f(x_t, y_t)$
end for

C.3 Approximate Generalized FTPL Lemma

Lemma C.5 (ϵ-FTPL Lemma)**.** *For the Generalized ϵ-FTPL algorithm,*

$$\text{REGRET} \leq \mathbb{E}\left[\sum_{t=1}^{T} f(x_{t+1}, y_t) - f(x_t, y_t)\right] + \mathbb{E}\left[\boldsymbol{\alpha} \cdot (\mathbf{\Gamma}_{x_1} - \mathbf{\Gamma}_{x^*})\right] + \epsilon T \tag{21}$$

where $x^* = \arg\max_{x \in \mathcal{X}} \sum_{t=1}^{T} f(x, y_t)$.

We first prove an approximate variant of Be-the-Leader Lemma.

Lemma C.6 (Be-the-Approximate-Leader Lemma)**.** *Under Generalized ϵ-FTPL, for any $x \in \mathcal{X}$,*

$$\sum_{t=1}^{T} f(x_{t+1}, y_t) + \boldsymbol{\alpha} \cdot \mathbf{\Gamma}_{x_1} \geq \sum_{t=1}^{T} f(x, y_t) + \boldsymbol{\alpha} \cdot \mathbf{\Gamma}_x - \epsilon T.$$

Proof. For $T = 1$ the inequality holds trivially, by the definition of approximate leader. Assume that the claim holds for some T. Then, for all x

$$\begin{aligned}
\sum_{t=1}^{T+1} f(x_{t+1}, y_t) + \boldsymbol{\alpha} \cdot \mathbf{\Gamma}_{x_1} &= \sum_{t=1}^{T} f(x_{t+1}, y_t) + \boldsymbol{\alpha} \cdot \mathbf{\Gamma}_{x_1} + f(x_{T+2}, y_{T+1}) \\
&\geq \sum_{t=1}^{T} f(x_{T+2}, y_t) + \boldsymbol{\alpha} \cdot \mathbf{\Gamma}_{x_{T+2}} - \epsilon T + f(x_{T+2}, y_{T+1}) \\
&\qquad \text{(by induction hypothesis)} \\
&= \sum_{t=1}^{T+1} f(x_{T+2}, y_t) + \boldsymbol{\alpha} \cdot \mathbf{\Gamma}_{x_{T+2}} - \epsilon T \\
&\geq \sum_{t=1}^{T+1} f(x, y_t) + \boldsymbol{\alpha} \cdot \mathbf{\Gamma}_x - \epsilon(T+1), \quad \text{(by approximate optimality of } x_{T+2})
\end{aligned}$$

proving the lemma. □

Proof of Lemma C.5. Let $x^* = \arg\max_{x \in \mathcal{X}} \sum_{t=1}^{T} f(x, y_t)$. Then by Be-the-Approximate-Leader Lemma,

$$\text{REGRET} = \mathbb{E}\left[\sum_{t=1}^{T} f(x^*, y_t) - \sum_{t=1}^{T} f(x_t, y_t)\right]$$

$$= \mathbb{E}\left[\sum_{t=1}^{T} f(x^*, y_t) - \sum_{t=1}^{T} f(x_{t+1}, y_t)\right] + \mathbb{E}\left[\sum_{t=1}^{T} f(x_{t+1}, y_t) - \sum_{t=1}^{T} f(x_t, y_t)\right]$$

$$\leq \mathbb{E}\left[\boldsymbol{\alpha} \cdot (\mathbf{\Gamma}_{x_1} - \mathbf{\Gamma}_{x^*})\right] + \mathbb{E}\left[\sum_{t=1}^{T} f(x_{t+1}, y_t) - \sum_{t=1}^{T} f(x_t, y_t)\right] + \epsilon T. \qquad \square$$

C.4 The Stability Lemma For Approximate Generalized FTPL

Lemma C.7. *Consider the Generalized ϵ-FTPL algorithm with a (κ, δ)-admissible matrix Γ with N columns and a $\left(\rho, \frac{1+2\epsilon}{\delta}\right)$-dispersed distribution D. Then,*

$$\mathbb{E}\left[\sum_{t=1}^{T} f(x_{t+1}, y_t) - f(x_t, y_t)\right] \leq 2TN\kappa\rho.$$

Proof. Fix any $t \leq T$. The bulk of the proof will establish that, with high probability, $\mathbf{\Gamma}_{x_{t+1}} = \mathbf{\Gamma}_{x_t}$, which by admissibility implies that $x_{t+1} = x_t$ and therefore $f(x_{t+1}, y_t) - f(x_t, y_t) = 0$.

Fix any $j \leq N$. We first show that $\Gamma_{x_{t+1}j} = \Gamma_{x_t j}$ with high probability. Let V denote the set of values that appear in the j^{th} column of $\mathbf{\Gamma}$. For any value $v \in V$, let x^v be the action that maximizes the perturbed profit among those whose $\mathbf{\Gamma}$ entry in the j^{th} column equals v:

$$x^v \coloneqq \operatorname*{arg\,max}_{x \in \mathcal{X}:\, \Gamma_{xj} = v} \left[\sum_{\tau=1}^{t-1} f(x, y_\tau) + \boldsymbol{\alpha} \cdot \mathbf{\Gamma}_x - \alpha_j v\right].$$

For any $v, v' \in V$, define

$$\Delta_{vv'} = \left(\sum_{\tau=1}^{t-1} f(x^v, y_\tau) + \boldsymbol{\alpha} \cdot \mathbf{\Gamma}_{x^v} - \alpha_j v\right) - \left(\sum_{\tau=1}^{t-1} f(x^{v'}, y_\tau) + \boldsymbol{\alpha} \cdot \mathbf{\Gamma}_{x^{v'}} - \alpha_j v'\right).$$

Note that x^v and $\Delta_{vv'}$ are independent of α_j, as we removed the payoff perturbation corresponding to α_j.

If $\Gamma_{x_t j} = v$, then by the ϵ-optimality of x_t on the perturbed profit, we have $\alpha_j(v' - v) - \epsilon \leq \Delta_{vv'}$ for all $v' \in V$. Suppose $\Gamma_{x_{t+1}j} \neq v$. Then there is some $v' \in V$ which yields a better perturbed profit, up to ϵ, in the next round, i.e.,

$$\sum_{\tau=1}^{t-1} f(x^{v'}, y_\tau) + f(x^{v'}, y_t) + \boldsymbol{\alpha} \cdot \mathbf{\Gamma}_{x^{v'}} \geq \sum_{\tau=1}^{t-1} f(x^v, y_\tau) + f(x^v, y_t) + \boldsymbol{\alpha} \cdot \mathbf{\Gamma}_{x^v} - \epsilon.$$

Rearranging, we obtain for this same v' that

$$\Delta_{vv'} \leq \alpha_j(v' - v) + f(x^{v'}, y_t) - f(x^v, y_t) + \epsilon \leq \alpha_j(v' - v) + 1 + \epsilon.$$

If $v' > v$, then

$$\alpha_j \geq \frac{\Delta_{vv'} - 1 - \epsilon}{v' - v} \geq \min_{\hat{v} \in V,\, \hat{v} > v} \frac{\Delta_{v\hat{v}} - 1 - \epsilon}{\hat{v} - v}$$

and so $\alpha_j(\overline{v} - v) + 1 + \epsilon \geq \Delta_{v\overline{v}}$ where $\overline{v}$ is the value of $\hat{v}$ minimizing the expression on the right. Thus, in this case we have $-\epsilon \leq \Delta_{v\overline{v}} - \alpha_j(\overline{v} - v) \leq 1 + \epsilon$. Similarly, if $v' < v$, then

$$\alpha_j \leq \frac{\Delta_{vv'} - 1 - \epsilon}{v' - v} \leq \max_{\hat{v} \in V,\, \hat{v} < v} \frac{\Delta_{v\hat{v}} - 1 - \epsilon}{\hat{v} - v}$$

and so $\alpha_j(\underline{v} - v) + 1 + \epsilon \geq \Delta_{v\underline{v}}$ where $\underline{v}$ is the value maximizing the expression on the right. In this case we have $-\epsilon \leq \Delta_{v\underline{v}} - \alpha_j(\underline{v} - v) \leq 1 + \epsilon$. Putting this all together, we have

$$\begin{aligned}
&\Pr\left[\Gamma_{x_{t+1}j} \neq \Gamma_{x_t j} \,\middle|\, \alpha_k, k \neq j\right] \\
&\quad \leq \Pr\left[\exists v \in V : \ -\epsilon \leq \Delta_{v\overline{v}} - \alpha_j(\overline{v} - v) \leq 1 + \epsilon \ \text{ or } \ -\epsilon \leq \Delta_{v\underline{v}} - \alpha_j(\underline{v} - v) \leq 1 + \epsilon \,\middle|\, \alpha_k, k \neq j\right] \\
&\quad \leq \sum_{v \in V} \left(\Pr\left[\alpha_j \in \left[\tfrac{\Delta_{v\overline{v}} - 1 - \epsilon}{\overline{v} - v}, \tfrac{\Delta_{v\overline{v}} + \epsilon}{\overline{v} - v}\right] \,\middle|\, \alpha_k, k \neq j\right] + \Pr\left[\alpha_j \in \left[\tfrac{-\Delta_{v\underline{v}} - \epsilon}{v - \underline{v}}, \tfrac{-\Delta_{v\underline{v}} + 1 + \epsilon}{v - \underline{v}}\right] \,\middle|\, \alpha_k, k \neq j\right] \right) \\
&\quad \leq 2\kappa\rho.
\end{aligned}$$

The last line follows from the fact that $\overline{v} - v \geq \delta$ and $v - \underline{v} \geq \delta$, the fact that D is $\left(\rho, \frac{1+2\epsilon}{\delta}\right)$-dispersed, and a union bound.

Since this bound does not depend on the values of the α_j, we can remove the conditioning and bound $\Pr[\Gamma_{x_{t+1}j} \neq \Gamma_{x_t j}] \leq 2\kappa\rho$. Taking a union bound over all $j \leq N$, we then have that, by admissibility, $\Pr[x_{t+1} \neq x_t] = \Pr\left[\mathbf{\Gamma}_{x_{t+1}} \neq \mathbf{\Gamma}_{x_t}\right] \leq 2N\kappa\rho$, which implies the result.

□

D Omitted Proofs From Applications to Online Auction Design

D.1 Proof of Equation (6)

Consider any vector of reserves $\mathbf{r} \in \mathcal{I}$ and let $\mathbf{r}' \in \mathcal{I}_m$ be the vector obtained by rounding each reserve price down to the nearest multiple of $1/m$. If $v_i > r_i$, then $v_i > r'_i$, so any bidder who would have been included in the basic VCG auction using reserves $\mathbf{r}$ is still included with $\mathbf{r}'$. This can only increase the number of bidders who are serviced and therefore pay a charge.

Now consider the amount that serviced bidder i is charged. This is the maximum of r_i and the highest bid of a bidder in the basic VCG auction who was not serviced (or 0 if all bidders were serviced); let b denote this highest unserviced bids in the basic VCG auction under $\mathbf{r}$, and similarly let b' denote such a bit under $\mathbf{r}'$. Since the set of bidders entering the basic VCG auction increases from $\mathbf{r}$ to $\mathbf{r}'$, we must have $b' \geq b$.

Let U be the set of bidders serviced under $\mathbf{r}$, and U' the set under $\mathbf{r}'$. The difference in revenue is

$$\begin{aligned}
&\sum_{i \in U} \max\{r_i, b\} - \sum_{i \in U'} \max\{r'_i, b'\} \\
&\qquad = \sum_{i \in U \cap U'} (\max\{r_i, b\} - \max\{r'_i, b'\}) + \sum_{i \in U \setminus U'} \max\{r_i, b\} - \sum_{i' \in U' \setminus U} \max\{r'_{i'}, b'\}. \qquad (22)
\end{aligned}$$

We begin by analyzing the last two terms. For any $i \in U \setminus U'$ and $i' \in U' \setminus U$,

$$r'_{i'} + 1/m > r_{i'} > v_{i'} \geq v_i > r_i,$$

where $v_{i'} \geq v_i$ follows, because i enters the basic VCG auction for $\mathbf{r}'$, but is not allocated the item. Therefore,

$$\max\{r'_{i'}, b'\} \geq \max\{r_i - 1/m,\, b\} \geq \max\{r_i, b\} - 1/m.$$

Since $|U \setminus U'| \leq |U' \setminus U|$, we can pick $V \subseteq U' \setminus U$ such that $|V| = |U \setminus U'|$ and obtain

$$\sum_{i \in U \setminus U'} \max\{r_i, b\} - \sum_{i' \in U' \setminus U} \max\{r'_{i'}, b'\} \leq \sum_{i \in U \setminus U'} \max\{r_i, b\} - \sum_{i' \in V} \max\{r'_{i'}, b'\} \leq \frac{|U \setminus U'|}{m}.$$

Note that each term in the first sum of Equation (22) is at most $1/m$, because

$$\max\{r'_i, b'\} \geq \max\{r_i - 1/m,\, b\} \geq \max\{r_i, b\} - 1/m.$$

Thus, we have

$$\mathrm{Rev}(\mathbf{r},\mathbf{v}) - \mathrm{Rev}(\mathbf{r}',\mathbf{v}) \le \frac{|U \cap U'|}{m} + \frac{|U \setminus U'|}{m} \le \frac{s}{m} \ .$$

This yields the approximation result

$$\max_{\mathbf{r}\in\mathcal{I}} \sum_{t=1}^{T} \mathrm{Rev}(\mathbf{r},\mathbf{v}_t) - \max_{\mathbf{r}\in\mathcal{I}_m} \sum_{t=1}^{T} \mathrm{Rev}(\mathbf{r},\mathbf{v}_t) \le \frac{Ts}{m}. \tag{23}$$

D.2 Omitted Part of the Proof of Lemma 4.2

We need to show that the sum of weights in any set $S_j(\alpha_j)$ is bounded by $\mathrm{poly}(n, m, \alpha_{\max})$. Recall that

$$w_m = \max\Big\{0,\ \max_k \big[\alpha_j m\big(k_\beta - (k-1)_\beta\big)\big]\Big\},$$

and for all $k = m, m-1, \ldots, 2$,

$$w_{k-1} = \frac{1}{k-1}\left(\sum_{h=k}^{m} w_h - \alpha_j m\big(k_\beta - (k-1)_\beta\big)\right).$$

By definition, $w_m \le m\alpha_{\max}$. We can also bound w_{k-1} as

$$w_{k-1} \le \frac{1}{k-1}\left(\sum_{h=k}^{m} w_h + m\alpha_{\max}\right) \le \frac{1}{k-1}\left(\sum_{h=k}^{m-1} w_h + 2m\alpha_{\max}\right),$$

so

$$(k-1)w_{k-1} \le \sum_{h=k}^{m-1} w_h + 2m\alpha_{\max}.$$

Summing over $k = 2, \ldots, m$, yields

$$\sum_{k=2}^{m}(k-1)w_{k-1} \le \sum_{h=2}^{m-1}(h-1)w_h + 2(m-1)m\alpha_{\max},$$

and subtracting the sum on the right-hand side yields

$$\sum_{k=2}^{m} w_{k-1} \le 2(m-1)m\alpha_{\max},$$

so $\sum_{k=1}^{m} w_k \le (2m^2 - 2m + m)\alpha_{\max}$.

D.3 Proof of Lemma 4.6

We will argue that the setting here is isomorphic to the setting in the proof of Lemma 4.2, so we can directly apply the result of analysis of $\mathbf{\Gamma}^{\mathrm{VCG}}$. The isomorphism from the VCG setting to IP setting maps bidders i in VCG to items ℓ in IP, and reserve price vectors $\mathbf{r}$ to price vectors $\mathbf{a}$. We therefore assume that n in VCG equals k in IP, and the values of m in both settings are equal. Then, indeed $\mathbf{\Gamma}^{\mathrm{VCG}}$ equals $\mathbf{\Gamma}^{\mathrm{IP}}$.

Next we need to show how to construct $S_j(\alpha_j)$ in the IP setting. Assume that j corresponds to the bidder i and the bit β in VCG setting, and the item ℓ and the bit β in IP setting. In VCG, we considered the bid profiles $\mathbf{v}_h = (h/m)\mathbf{e}_i$, and the revenue of any auction $\mathbf{r}$ is

$$\mathrm{Rev}^{\mathrm{VCG}}(\mathbf{r},\mathbf{v}_h) = r_i \mathbf{1}_{(h \ge m r_i)}.$$

In IP setting, we consider profiles $\mathbf{v}'_h$ of combinatorial valuations over bundles $\mathbf{q} \in \{0,1\}^k$, in which all bidders have values zero on all bundles and one bidder has value h/m for bundle $\mathbf{e}_\ell$ and zero on all other bundles.[6] In this case, we have

$$\mathrm{Rev}^{\mathrm{IP}}(\mathbf{a}, \mathbf{v}'_h) = a_i \mathbf{1}_{(h \geq m a_i)}.$$

Thus, both the translation matrices $\mathbf{\Gamma}^{\mathrm{VCG}}$ and $\mathbf{\Gamma}^{\mathrm{IP}}$ as well as the revenue functions $\mathrm{Rev}^{\mathrm{VCG}}$ and $\mathrm{Rev}^{\mathrm{IP}}$ are isomorphic (given these choices of the profiles). Therefore, we can set the weights w'_h in IP setting equal to the weights w_h in VCG setting and obtain admissibility and implementability with the same constants and complexity.

D.4 Proof of Lemma 4.12

Since $\Gamma^{\mathrm{IP}}_{\boldsymbol{\theta},\mathbf{v}} = \mathrm{Rev}(\boldsymbol{\theta}, \mathbf{v})$, for any $\alpha \geq 0$, $S_{\mathbf{v}}(\alpha) = \{(\alpha_{\mathbf{v}}, \mathbf{v})\}$. Furthermore, the mass of dataset $S_{\mathbf{v}}(\alpha)$ is $\mathrm{poly}(n, m, \alpha_{\max})$. So, $(\mathbf{\Gamma}, D)$ is implementable with complexity $\mathrm{poly}(n, m, \alpha_{max})$.

Take any two different auctions $\boldsymbol{\theta}$ and $\boldsymbol{\theta}'$. We show that $\mathbf{\Gamma}^{\mathrm{IP}}_{\boldsymbol{\theta}} \neq \mathbf{\Gamma}^{\mathrm{IP}}_{\boldsymbol{\theta}'}$. Let b be the smallest level at which there is $i \in [n]$ such that $\theta^i_b \neq \theta'^i_b$ and among such i choose the largest. There are three cases:

1. $i \neq n$: Consider the bid profile $\mathbf{v}^{i,\ell}$ for $\ell = m\theta^n_b$. By the choice of i and the fact that $i \neq n$, we have that $b^{\boldsymbol{\theta}}_n(v^{i,\ell}_n) = b^{\boldsymbol{\theta}'}_n(v^{i,\ell}_n) = b$. On the other hand, $b^{\boldsymbol{\theta}}_i(v^{i,\ell}_i) = s - 1 \geq b$. Therefore, bidder i wins the item in both auctions and pays the b^{th} threshold. So, $\mathrm{Rev}(\boldsymbol{\theta}, \mathbf{v}^{i,\ell}) = \theta^i_b \neq \theta'^i_b = \mathrm{Rev}(\boldsymbol{\theta}', \mathbf{v}^{i,\ell})$.
2. $i = n$ and $b \geq 1$: WLOG, assume that $\theta^n_b < \theta'^n_b$. Let $\ell = m\theta^n_b$ and consider $\mathbf{v}^{1,\ell}$. Then $b^{\boldsymbol{\theta}}_n(v^{1,\ell}_n) = b$ and $b^{\boldsymbol{\theta}'}_n(v^{1,\ell}_n) = b'$ for some $b' < b$. So, bidder 1 wins the item in both auctions and pays the threshold that corresponds to the n^{th} bidder's level. Therefore, $\mathrm{Rev}(\boldsymbol{\theta}, \mathbf{v}^{1,\ell}) = \theta^1_b \neq \theta^1_{b'} = \mathrm{Rev}(\boldsymbol{\theta}', \mathbf{v}^{1,\ell})$.
3. $i = n$ and $b = 0$: Consider bid profile $\mathbf{e}_n$. In this profile, bidder n wins and pays the absolute reserve price. Therefore, $\mathrm{Rev}(\boldsymbol{\theta}, \mathbf{e}_n) = \theta^n_0 \neq \theta'^n_0 = \mathrm{Rev}(\boldsymbol{\theta}', \mathbf{e}_n)$.

Therefore, $\mathbf{\Gamma}^{\mathrm{IP}}_{\boldsymbol{\theta}} \neq \mathbf{\Gamma}^{\mathrm{IP}}_{\boldsymbol{\theta}'}$. Since any element of $\mathbf{\Gamma}^{\mathrm{IP}}$ is a multiple of $1/m$ then $\mathbf{\Gamma}^{\mathrm{IP}}$ is $(m+1, \frac{1}{m})$-admissible.

D.5 Proof of Lemma 4.14

We say that two auction $\boldsymbol{\theta}$ and $\boldsymbol{\theta}'$ are *equivalent* if for all $\mathbf{v}$, $\mathrm{Rev}(\boldsymbol{\theta}, \mathbf{v}) = \mathrm{Rev}(\boldsymbol{\theta}', \mathbf{v})$. In what follows, for any $\boldsymbol{\theta} \in \mathcal{R}_{s,m}$ we construct $\boldsymbol{\theta}' \in \overline{\mathcal{R}}_{s,m}$, such that $\boldsymbol{\theta}$ and $\boldsymbol{\theta}'$ are equivalent.

When $\boldsymbol{\theta} \in \mathcal{S}_{s,m}$, then $\boldsymbol{\theta}$ has no empty buckets. That is, for all i and b, $\theta^i_b \in B^i_b(\boldsymbol{\theta})$. So, the claim holds trivially for $\boldsymbol{\theta}' = \boldsymbol{\theta}$.

Let $\boldsymbol{\theta}$ be an auction such that for some i and b, $B^i_b(\boldsymbol{\theta}) \neq \emptyset$ but all of its competing buckets are empty. Consider auction $\boldsymbol{\theta}'$ with thresholds $\theta'^i_b = \theta^i_{b+1}$ and $\theta'^{i'}_{b'} = \theta^{i'}_{b'}$ for all $i' \neq i$ or $b' \neq b$. We show that $\boldsymbol{\theta}$ and $\boldsymbol{\theta}'$ are equivalent. Consider any bid profile $\mathbf{v}$. There are two cases: 1) $v_i \notin B^i_b(\boldsymbol{\theta})$; 2) $v_i \in B^i_b(\boldsymbol{\theta})$.

Case 1: If $v_i \notin B^i_b(\boldsymbol{\theta})$, then it is easy to see that $b^{\boldsymbol{\theta}}_i(v_i) = b^{\boldsymbol{\theta}'}_i(v_i)$ so the winner remains the same. As a result the payment would remain the same if i is not the winner. When i is the winner, since all of the competing buckets of $B^i_b(\boldsymbol{\theta})$ are empty, the payment of bidder i is not at level b, so the payment of bidder i also remains the same. Therefore, $\mathrm{Rev}(\boldsymbol{\theta}, \mathbf{v}) = \mathrm{Rev}(\boldsymbol{\theta}', \mathbf{v})$.

Case 2: If $v_i \in B^i_b$, then $b^{\boldsymbol{\theta}'}_i(v_i) = b^{\boldsymbol{\theta}}_i(v_i) - 1$. Let i' be any bidder that appears after bidder i in the sorted list of the bid level. Since all competing buckets of B^i_b are empty, it must be that either $i < i'$ and $b^{\boldsymbol{\theta}}_{i'}(v_{i'}) < b$ or $i > i'$ and $b^{\boldsymbol{\theta}}_i(v_{i'}) < b - 1$. In either case, position of i in the sorted list remains the same in $\boldsymbol{\theta}'$. Consider the following three cases:

1. Bidder i is the winner: Let i' be the runner up bidder. As described above, i is still the first in the decreasing order of bid levels. So, i wins the item in $\boldsymbol{\theta}'$. Moreover, because the competing buckets of

[6] Note that a simple variation of this bid profile can be used in settings where the valuations need to satisfy additional assumptions, such as (sub-)additivity or free disposal. In such cases, we can use a similar bid profile where one bidder has valuation h/m for any bundle that includes item ℓ and all other valuations are 0.

$B_b^i(\boldsymbol{\theta})$ are empty, the payment of bidder i in $\boldsymbol{\theta}$ is at a threshold with level below $\leq b-1$, which is the same in both auctions. Therefore, $\text{Rev}(\boldsymbol{\theta}, \mathbf{v}) = \text{Rev}(\boldsymbol{\theta}', \mathbf{v})$.

2. $i' \neq i$ wins the item and i is the runner up: As described above, i is still the runner up in the decreasing order of bid levels in $\boldsymbol{\theta}'$. There are two cases:
   (a) If $i' < i$, then $\text{Rev}(\boldsymbol{\theta}, \mathbf{v}) = \theta_b^{i'}$. Note that $B_{b-1}^{i'} = \emptyset$, so, $\theta_{b-1}^{i'} = \theta_b^{i'}$. So, $\text{Rev}(\boldsymbol{\theta}, \mathbf{v}) = \theta_b^{i'} = \theta_{b-1}^{i'} = \text{Rev}(\boldsymbol{\theta}', \mathbf{v})$.
   (b) If $i' > i$, then $\text{Rev}(\boldsymbol{\theta}, \mathbf{v}) = \theta_{b+1}^{i'}$. Note that $B_b^{i'} = \emptyset$, so, $\theta_b^{i'} = \theta_{b+1}^{i'}$. So, $\text{Rev}(\boldsymbol{\theta}, \mathbf{v}) = \theta_{b+1}^{i'} = \theta_b^{i'} = \text{Rev}(\boldsymbol{\theta}', \mathbf{v})$.
   So, the revenue remains the same.
3. If i is not the winner or the runner up: As described above i is not the winner or the runner up in $\boldsymbol{\theta}'$, either. So, the winner and the revenue are not affected.

Therefore, $\text{Rev}(\boldsymbol{\theta}, \mathbf{v}) = \text{Rev}(\boldsymbol{\theta}', \mathbf{v})$ for all $\mathbf{v}$. Furthermore, the thresholds in $\boldsymbol{\theta}'$ are all equal or larger than the corresponding thresholds in $\boldsymbol{\theta}$, with one threshold being strictly larger. So repeating the above process for other buckets that violate the claim of the lemma will eventually halt. This proves the lemma.

D.6 Proof of Lemma 4.15

In the interest of readability, we drop the superscript and write $\boldsymbol{\Gamma}$ for $\boldsymbol{\Gamma}^{\text{RL}}$ in this proof.

Since $\Gamma_{\boldsymbol{\theta},\mathbf{v}} = \text{Rev}(\boldsymbol{\theta}, \mathbf{v})$, it is immediately clear that $(\boldsymbol{\Gamma}, D)$ can be implemented by dataset $S_{\mathbf{v}}(\alpha_{\mathbf{v}}) = \{(\alpha_{\mathbf{v}}, \mathbf{v})\}$. Furthermore, the mass of $S_{\mathbf{v}}(\alpha_{\mathbf{v}})$ is $\text{poly}(n, m, \alpha_{\max})$. So, $(\boldsymbol{\Gamma}, D)$ is implementable with complexity $\text{poly}(n, m, \alpha_{max})$.

Next, we show that for any two $\boldsymbol{\theta}$ and $\boldsymbol{\theta}'$, $\boldsymbol{\Gamma}_{\boldsymbol{\theta}} \neq \boldsymbol{\Gamma}_{\boldsymbol{\theta}'}$. Let b be the smallest integer for which there exists $i \in [n]$, such that $\theta_b^i \neq \theta_b'^i$. Choose i as the largest bidder for which $\theta_b^i \neq \theta_b'^i$. If $b = 0$: Then $\text{Rev}(\boldsymbol{\theta}, \mathbf{e}_i) = \theta_0^i \neq \theta_0'^i = \text{Rev}(\boldsymbol{\theta}', \mathbf{e}_i)$. So, $\boldsymbol{\Gamma}_{\boldsymbol{\theta}} \neq \boldsymbol{\Gamma}_{\boldsymbol{\theta}'}$.

If $b > 0$, then the $(b-1)^{\text{th}}$ bucket is non-empty in one of the auctions — WLOG, let that be auction $\boldsymbol{\theta}$. This implies that $\theta_b^i > \theta_b'^i$. By the structure of $\overline{\mathcal{R}}_{s,m}$, there is a non-empty competing bucket.

If the non-empty competing bucket is $B_{b-1}^{i'}(\boldsymbol{\theta})$ for $i' > i$, then consider the bid profile $\mathbf{v}^{i',i,\ell}$ for $\ell = m\theta_b'^i$. Since $B_{b-1}^{i'}(\boldsymbol{\theta}) \neq \emptyset$ and i was the largest bidder for which the thresholds differ, we have that $B_{b-1}^{i'}(\boldsymbol{\theta}') \neq \emptyset$ and $\theta_b^i = \theta_b'^i \neq \theta_{b-1}^i = \theta_{b-1}^i$. Furthermore, $b_i^{\boldsymbol{\theta}}(v_i^{i',i,\ell}) < b$ and $b_i^{\boldsymbol{\theta}'}(v_i^{i',i,\ell}) \geq b$. So, i' wins the item and pays the level of bidder $< b$ in $\boldsymbol{\theta}$ and $\geq b$ in $\boldsymbol{\theta}'$, which are different. So, $\text{Rev}(\boldsymbol{\theta}, \mathbf{v}^{i',i,\ell}) \neq \text{Rev}(\boldsymbol{\theta}', \mathbf{v}^{i',i,\ell})$. The case where the non-empty competing bucket is $B_{b-2}^{i'}(\boldsymbol{\theta})$ for $i' < i$ is similar.

Therefore, $\boldsymbol{\Gamma}_{\boldsymbol{\theta}} \neq \boldsymbol{\Gamma}_{\boldsymbol{\theta}'}$. Since any element of $\boldsymbol{\Gamma}$ is a multiple of $1/m$, then $\boldsymbol{\Gamma}$ is $(m+1, \frac{1}{m})$-admissible.

E Proof of Equation 15

We use ideas similar to the ones discussed by Morgenstern and Roughgarden [28] and Devanur et al. [12] to show that $\mathcal{R}_{s,m}$, the class of s-level auctions on a discrete grid, approximates the revenue of Myerson auction within an additive value of $O(\epsilon)$ for appropriate parameters $m = O(1/\epsilon)$ and $s = O(1/\epsilon)$.

Let $m = 1/(2\epsilon)$ and $s = 1/\epsilon$.

First, let F' be a product distribution obtained by rounding down each $\mathbf{v} \sim F$ to the closest *odd* multiple of ϵ. We show that $\textsc{Opt}(F') \geq \textsc{Opt}(F) - 2\epsilon$. Let M be the optimal Myerson auction for F and let M' be the following mechanism for allocating items to $\mathbf{v}' \sim F'$: Take $\mathbf{v}$ such that $1 - F_i(v_i) = 1 - F_i'(v_i')$ for all $i \in [n]$ and allocate the item according to the outcome of mechanism M on $\mathbf{v}$. Charge the winner the minimum value above which it would remain a winner.

To analyze the revenue of M', consider any vector of quantiles $\mathbf{q} \in [0,1]^n$ and let $\mathbf{v}$ and $\mathbf{v}'$ be the values in F and F' corresponding to these quantiles, i.e., $1 - F_i(v_i) = 1 - F_i'(v_i') = q_i$ for all $i \in [n]$. Note that, allocation of M on $\mathbf{v}$ is the same as the allocation of M' on $\mathbf{v}'$. Moreover, the quantiles of the payments are the same in both mechanisms. So, for any such quantile the payment in F and F' differ by at most 2ϵ. This prove that $\textsc{Opt}(F') \geq \textsc{Opt}(F) - 2\epsilon$.

Next, we show that there is $\boldsymbol{\theta} \in \mathcal{R}_{1/\epsilon,\infty}$ such that $\mathbb{E}_{\mathbf{v}\sim F'}[\mathrm{Rev}(\boldsymbol{\theta}, \mathbf{v})] \geq \textsc{Opt}(F') - \epsilon$. For each i, let $\phi'_i(\cdot)$ be the ironed virtual valuation function for bidder i with respect to F'_i. These ϕ'_i are monotone functions and can be made into strictly increasing functions by removing an arbitrary small fraction of the distribution and losing an arbitrarily small fraction of the revenue. Let $\boldsymbol{\theta}$ be such that $\theta^i_b = \phi^{-1}(b\epsilon)$ for $b \in \{0, 1, \ldots, s-1\}$. Next, we show that for any $\mathbf{v}$, $\mathrm{Rev}(\boldsymbol{\theta}, \mathbf{v}) \geq \textsc{Opt}(F') - \epsilon$.

Consider $\mathbf{v} \sim F'$ and let i^* and i' be the winners in $\boldsymbol{\theta}$ and the Myerson optimal auctions respectively. Since in s-threshold and Myerson optimal auctions the allocations only depend on the ironed virtual values and not the values themselves, we have that[7]

$$\mathbb{E}_{\mathbf{v}\sim F'}[\mathrm{Rev}(\boldsymbol{\theta}, \mathbf{v})] = \mathbb{E}_{\mathbf{v}\sim F'}[\phi_{i^*}(v_{i^*})] \quad \text{and} \quad \textsc{Opt}(F') = \mathbb{E}_{\mathbf{v}\sim F'}[\phi_{i'}(v_{i'})].$$

Now, consider the winners i^* and i'. Note that if i^* was the unique bidder at the highest bucket under $\boldsymbol{\theta}$ (there were no ties to be broken lexicographically), then i^* also has the highest virtual valuation, so $i^* = i'$. On the other hand, if i' was tied with i^*, then $\phi_{i'}(v_{i'}) - \phi_{i^*}(v_{i^*}) \leq \epsilon$. So, overall we have

$$\mathbb{E}_{\mathbf{v}\sim F'}[\mathrm{Rev}(\boldsymbol{\theta}, \mathbf{v})] = \mathbb{E}_{\mathbf{v}\sim F'}[\phi_{i^*}(v_{i^*})] \geq \mathbb{E}_{\mathbf{v}\sim F'}[\phi_{i'}(v_{i'})] - \epsilon = \textsc{Opt}(F') - \epsilon.$$

Now we show how to round the thresholds of $\boldsymbol{\theta} \in \mathcal{R}_{1/\epsilon,\infty}$ to $\boldsymbol{\theta}' \in \mathcal{R}_{1/\epsilon,1/\epsilon}$, without losing a large fraction of the revenue. For any $\boldsymbol{\theta}$, and for any $\theta_i \in [(2j-1)\epsilon, (2j+1)\epsilon)$, let $\theta'_i = 2j\epsilon$. By construction F' is only supported on odd multiples of ϵ. So, the allocation remains the same for any $\mathbf{v} \in \mathrm{supp}(F')$ under $\boldsymbol{\theta}$ and $\boldsymbol{\theta}'$. And the payment reduces by at most ϵ. So,

$$\mathbb{E}_{\mathbf{v}\sim F'}[\mathrm{Rev}(\boldsymbol{\theta}', \mathbf{v})] \geq \mathbb{E}_{\mathbf{v}\sim F'}[\mathrm{Rev}(\boldsymbol{\theta}, \mathbf{v})].$$

Moreover, since each bid profiles in F maps to bid profile with smaller bids in F', we

$$\mathbb{E}_{\mathbf{v}\sim F}[\mathrm{Rev}(\boldsymbol{\theta}', \mathbf{v})] \geq \mathbb{E}_{\mathbf{v}\sim F'}[\mathrm{Rev}(\boldsymbol{\theta}', \mathbf{v})] \geq \mathbb{E}_{\mathbf{v}\sim F'}[\mathrm{Rev}(\boldsymbol{\theta}, \mathbf{v})] - \epsilon \geq \textsc{Opt}(F') - 2\epsilon \geq \textsc{Opt}(F) - 4\epsilon.$$

This proves the claim.

F Proof of Theorem 6.8

Using Lemma 6.7, we have

$$\mathbb{E}\left[\sum_{t=1}^{T} f(x_t, y_t) - f(x_{t+1}, y_t)\right] \leq 2TN\kappa \cdot \max_{s\in\mathbb{R}} \Pr\left[Z \in [s, s + \tfrac{1}{\delta}]\right] \leq 2TN\kappa \frac{1}{2\delta\nu}$$

where $Z \sim U[-\nu, \nu]$. It remains to bound the term $\mathbb{E}\left[\sum_{\sigma\in S}\sum_{j=1}^{N} \alpha_{\sigma,j}(\Gamma^S_{x^*,(\sigma,j)} - \Gamma^S_{x_1,(\sigma,j)})\right]$, which is at most

$$2\mathbb{E}\left[\max_x \sum_{\sigma\in S}\sum_{j=1}^{N} \alpha_{\sigma,j}\Gamma_{x,(\sigma,j)}\right].$$

Let $\beta_x = \sum_{\sigma\in S}\sum_{j=1}^{N} \alpha_{\sigma,j}\Gamma_{x,(\sigma,j)}$. We therefore have for any $\lambda > 0$,

$$\begin{aligned}
\mathbb{E}\left[\max_x \beta_x\right] &= \tfrac{1}{\lambda}\ln\left(\exp\left(\lambda\mathbb{E}\left[\max_x \beta_x\right]\right)\right) \\
&\leq \tfrac{1}{\lambda}\ln\left(\mathbb{E}\left[\exp\left(\lambda \max_x \beta_x\right)\right]\right) && \text{(Jensen's inequality)} \\
&\leq \tfrac{1}{\lambda}\ln\left(\sum_x \mathbb{E}\left[\exp\left(\lambda\beta_x\right)\right]\right)
\end{aligned}$$

[7] See further discussion on this by Hartline [19] and Morgenstern and Roughgarden [28].

$$
\begin{aligned}
&\leq \tfrac{1}{\lambda} \ln \left(\sum_{x} \prod_{\sigma,j} \mathbb{E}\left[\exp\left(\lambda \alpha_{\sigma,j} \Gamma_{x,(\sigma,j)} \right) \right] \right) \\
&\leq \tfrac{1}{\lambda} \ln \left(\sum_{x} \prod_{\sigma,j} \exp\left(\tfrac{\nu^2 \gamma^2 \lambda^2}{2} \right) \right) && \text{(Hoeffding's lemma for bounded r.v.)} \\
&= \frac{\ln |\mathcal{X}|}{\lambda} + \frac{\nu^2 \gamma^2 N |S| \lambda}{2} \\
&= \sqrt{2N|S| \ln |\mathcal{X}|} \nu \gamma. && \text{(by picking the optimal } \lambda\text{)}
\end{aligned}
$$

Combining the stability and the error bound above and invoking Lemma 3.1, completes the proof.